

**COMMUNICATION SYSTEM  
WITH MULTICARRIER TELEPHONY TRANSPORT**

**Cross Reference to Related Cases**

5  
Sub  
al

10 This application is a continuation-in-part of U.S. Application Serial Nos. \_\_\_\_\_ (Atty Docket #500.615US1), 08/311964, 08/457295, and 08/457317, which applications are incorporated by reference. This application is related to U.S. Application Serial Nos. 08/384659, 08/455340, 08/455059, 08/457294, 08/457110, 08/456871, 08/457022, and 08/457037, which applications are incorporated by reference.

**Field of the Invention**

15 The present invention relates generally to the field of communication systems. More particularly, the present invention relates to communication systems with multicarrier telephony transport.

**Background of the Invention**

20 Two information services found in households and businesses today include television, or video, services and telephone services. Another information service involves digital data transfer which is most frequently accomplished using a modem connected to a telephone service. All further references to telephony herein shall include both telephone services and digital data transfer services.

25 Characteristics of telephony and video signals are different and therefore telephony and video networks are designed differently as well. For example, telephony information occupies a relatively narrow band when compared to the bandwidth for video signals. In addition, telephony signals are low frequency whereas NTSC standard video signals are transmitted at carrier frequencies greater than 50 MHz. Accordingly, telephone transmission networks are relatively narrow  
30 band systems which operate at audio frequencies and which typically serve the

customer by twisted wire drops from a curb-side junction box. On the other hand, cable television services are broad band and incorporate various frequency carrier mixing methods to achieve signals compatible with conventional very high frequency television receivers. Cable television systems or video services are typically provided by cable television companies through a shielded cable service connection to each individual home or business.

One attempt to combine telephony and video services into a single network is described in U.S. Patent No. 4,977,593 to Balance entitled "Optical Communications Network." Balance describes a passive optical communications network with an optical source located in a central station. The optical source transmits time division multiplexed optical signals along an optical fiber and which signals are later split by a series of splitters between several individual fibers servicing outstations. The network allows for digital speech data to be transmitted from the outstations to the central station via the same optical path. In addition, Balance indicates that additional wavelengths could be utilized to add services, such as cable television, via digital multiplex to the network.

A 1988 NCTA technical paper, entitled "Fiber Backbone: A Proposal For an Evolutionary Cable TV network Architecture," by James A. Chiddix and David M. Pangrac, describes a hybrid optical fiber/coaxial cable television (CATV) system architecture. The architecture builds upon existing coaxial CATV networks. The architecture includes the use of a direct optical fiber path from a head end to a number of feed points in an already existing CATV distribution system.

U.S. Patent No. 5,153,763 to Pidgeon, entitled "CATV Distribution Networks Using Light Wave Transmission Lines," describes a CATV network for distribution of broad band, multichannel CATV signals from a head end to a plurality of subscribers. Electrical to optical transmitters at the head end and optical to electrical receivers at a fiber node launch and receive optical signals corresponding to broad band CATV electrical signals. Distribution from the fiber node is obtained by transmitting electrical signals along coaxial cable transmission

lines. The system reduces distortion of the transmitted broad band CATV signals by block conversion of all or part of the broad band of CATV signals to a frequency range which is less than an octave. Related U.S. Patent No. 5,262,883 to Pidgeon, entitled "CATV Distribution Networks Using Light Wave Transmission Lines," further describes the distortion reducing system.

Although the above-mentioned networks describe various concepts for transmitting broad band video signals over various architectures, which may include hybrid optical fiber/coax architectures, none of these references describe a cost effective, flexible, communications system for telephony communications. Several problems are inherent in such a communication system.

One such problem is the need to optimize the bandwidth used for transporting data so that the bandwidth used does not exceed the allotted bandwidth. Bandwidth requirements are particularly critical in multi-point to point communication where multiple transmitters at remote units must be accommodated such that allotted bandwidth is not exceeded.

A second problem involves power consumption of the system. The communication system should minimize the power used at the remote units for the transport of data, as the equipment utilized at the remote units for transmission and reception may be supplied by power distributed over the transmission medium of the system.

Another problem arises from a fault in the system preventing communication between a head end and multiple remote units of a multi-point to point system. For example, a cut transmission line from a head end to many remote units may leave many users without service. After the fault is corrected, it is important bring as many remote units back into service as quickly as possible.

Data integrity must also be addressed. Both internal and external interference can degrade the communication. Internal interference exists between data signals being transported over the system. That is, transported data signals over a common communication link may experience interference therebetween,

decreasing the integrity of the data. Ingress from external sources can also effect the integrity of data transmissions. A telephony communication network is susceptible to "noise" generated by external sources, such as HAM radio. Because such noise can be intermittent and vary in intensity, a method of transporting data over the system should correct or avoid the presence of such ingress.

These problems and others as will become apparent from the description to follow, present a need for an enhanced communication system. Moreover, once the enhanced system is described, a number of practical problems in its physical realization are presented and overcome.

Another embodiment provides a method and apparatus for a fast Fourier transform. This invention relates to the field of electronic communication systems, and more specifically to an improved method and apparatus for providing a fast Fourier transform ("FFT").

There are many advanced digital signal-processing applications requiring analysis of large quantities of data in short time periods, especially where there is interest in providing "real time" results. Such applications include signal processing in modems which use OFDM (orthogonal frequency division multiplexing). In order to be useful in these and other applications, Discrete Fourier Transform (DFT) or Fast Fourier Transform (FFT) signal processors must accommodate large numbers of transforms, or amounts of data, in very short processing times, often called high data throughput.

In addition to the speed and data-throughput requirements, power consumption is a major concern for many applications. In some signal-processing applications, power is supplied by portable generation or storage equipment, such as batteries, where the ultimate power available is limited by many environment. In such applications, processor power consumption must be as low as possible. One useful measure of utility or merit for FFT processors is the energy dissipation per transform point. Ultimately, one key problem with any FFT processor is the amount of power consumed per transform. Generally, high-performance, efficient FFT

processors exhibit energy dissipations per transform in the range of 100 to 1000 times  $\log_2 N$  nanojoules, where  $N$  is the number of points in a given transform. As a consequence, reasonably large transforms required to process large arrays of data, result in large power consumption.

Machine-implemented computation of an FFT is often simplified by cascading together a series of simple multiply-and-add stages. When a recursive process is used, data circulates through a single stage and the computational structure of the stage is made variable for each circulation. Each circulation through the stage is referred to as a "pass".

A plurality of computational elements, each known as a radix- $r$  butterfly, may be assembled to define a single stage for carrying out a particular pass. A radix- $r$  butterfly receives  $r$  input signals and produces a corresponding number of  $r$  output signals, where each output signal is the weighted sum of the  $r$  input signals. The radix number,  $r$ , in essence, defines the number of input components which contribute to each output component.

By way of example, a radix-2 butterfly receives two input signals and produces two output signals. Each output signal is the weighted sum of the two input signals. A radix-4 butterfly receives four input signals and produces four corresponding output signals. Each output signal of the radix-4 butterfly constitutes a weighted sum of the four input signals.

Completion of an  $N$ -point Fast Fourier Transform (FFT) requires that the product of the butterfly radix values, taken over the total number of stages or passes, equals the total point count,  $N$ . Thus, a 64-point FFT can be performed by one radix-64 butterfly, or three cascaded stages where each stage has sixteen radix-4 butterflies (the product of the radix values for stage-1 and stage-2 and stage-3 is  $4 \times 4 \times 4 = 64$ ), or six cascaded stages where each of the six stages comprises 32 radix-2 butterflies (the product of the radix values for stage-1 through stage-6 is  $2 \times 2 \times 2 \times 2 \times 2 \times 2 = 64$ ).

A multi-stage or multi-pass FFT process can be correctly carried out under

2025 RELEASE UNDER E.O. 14176

conditions where the number of butterfly elements changes from one pass (or stage) to the next and the radix value,  $r$ , of the butterfly elements also changes from one pass (or stage) to the next. A paper by Gordon DeMuth, "ALGORITHMS FOR DEFINING MIXED RADIX FFT FLOW GRAPHS", IEEE Transactions on Acoustics, Speech, and Signal Processing, Vol 37, No. 9, September 1989, Pages 1349-1358, describes a generalized method for performing an FFT with a mixed-radix system. A mixed-radix system is one where the radix value,  $r$ , in one stage or pass is different from that of at least one other stage or pass.

An advantage of a mixed-radix computing system is that it can be "tuned" to optimize the signal-to-noise ratio of the transform (or more correctly speaking, to minimize the accumulated round-off error of the total transform) for each particular set of circumstances. By way of example, it is advantageous in one environment to perform a 512-point FFT using the mixed-radix sequence: 4, 4, 4, 4, 2. In a different environment, it may be more advantageous to use the mixed-radix sequence: 4, 2, 4, 4, 4. Round-off error varies within a machine of finite precision as a function of radix value and the peak signal magnitudes that develop in each stage or pass.

In addition, it may be advantageous to scale intermediate results between each stage or pass, in order to minimize round-off errors and the problem of overflow. Further, it may be advantageous to vary the amount of scaling performed between each pass, e.g., either to scale by  $1/4$  between each radix-4 stage or to scale by  $1/2$  for some stages and  $1/8$  for other stages.

Heretofore, FFT processors generally fetched data values from their working storage in a serial manner, thus limiting the speed which could be obtained. Further, current FFT processors generally were limited in speed by loading the working storage with input values, then processing the data in the working storage, then unloading the result values.

There are many advanced digital signal-processing applications requiring analysis of large quantities of data in short time periods, especially where there is

FOUO 44-38860

interest in providing "real time" results. Such applications include signal processing in modems which use OFDM (orthogonal frequency division multiplexing).

One need in the art is for an accurate analog-to-digital conversion (ADC) at moderate frequencies having limited bandwidth. One technology known in the art is the "Sigma-Delta" ADC which provides very good resolution (high number of bits in the digital result), but only for signals whose converted signal bandwidth is low.

Another need is for an ADC which provides bandwidth-limited digital I and Q signals (representing amplitude and quadrature) for a 200 kHz bandwidth received analog modem signal, wherein the digital result has very high resolution and accuracy.

What is needed is a method and apparatus which addresses the above problems in the art.

#### Summary of the Invention

The present invention describes a multi-point to point communication system including multicarrier telephony transport. The multi-point to point communication system includes a hybrid fiber/coax distribution network. A head end terminal provides for downstream transmission of downstream control data and downstream telephony information in a first frequency bandwidth over the hybrid fiber/coax distribution network and reception of upstream telephony information and upstream control data in a second frequency bandwidth over the hybrid fiber/coax distribution network. The head end terminal includes a head end multicarrier modem for modulating at least downstream telephony information on a plurality of orthogonal carriers in the first frequency bandwidth and demodulating at least upstream telephony information modulated on a plurality of orthogonal carriers in the second frequency bandwidth. The head end terminal further includes a head end controller operatively connected to the head end multicarrier modem for controlling transmission of the downstream telephony information and downstream control data and for controlling receipt of the upstream control data and upstream

telephony information. The system further includes at least one service unit, each service unit associated with at least one remote unit and operatively connected to the hybrid fiber/coax distribution network for upstream transmission of upstream telephony information and upstream control data in the second frequency bandwidth and for receipt of the downstream control data and downstream telephony information in the first frequency bandwidth. Each service unit includes a service unit multicarrier modem for modulating at least the upstream telephony information on at least one carrier orthogonal at the head end to at least one other carrier in the second frequency bandwidth and for demodulating at least the downstream telephony information modulated on at least a band of a plurality of orthogonal carriers in the first frequency bandwidth. Each service unit also includes a service unit controller operatively connected to the service unit multicarrier modem for controlling the modulation of and demodulation performed by the service unit multicarrier modem.

Another embodiment prevents untoward spectral effects in the multicarrier signal from variations in channel activity and from highly repetitive data patterns in the payload channels. Data in the payload channels can be scrambled with pseudorandom sequences, and different sequences can be applied to different channels in order to produce a more balanced multicarrier spectrum.

In another embodiment, the plurality of orthogonal carriers in the first frequency bandwidth include at least one control channel for transmission of downstream control data and a plurality of telephony information channels for transmission of downstream telephony information. Further, the plurality of orthogonal carriers in the second frequency bandwidth include at least one control channel for transmission of upstream control data and a plurality of telephony information channels for transmission of upstream telephony information.

In other embodiments, a plurality of control channels are interspersed among the telephony information channels in the first frequency bandwidth and a plurality of control channels are interspersed among the telephony channels of the second

10020" 42E70660



frequency bandwidth. The telephony channels may be divided into subbands each having multiple data or payload channels and a control channel; this allows the remote modems to be realized as less expensive, and/or better performing narrow-band modems.

5           The clock signals for generating the carriers and the symbols representing the transmitted data may be locked to each other or generated from the same source, to reduce intersymbol interference significantly.

10           Another technique for reducing intersymbol interference is the transmission of each symbol with more than  $360^\circ$  of phase in one cycle of its carrier, in order to allow some leeway in tracking the phase of a channel carrier in a receiving modem.

15           Some applications demand more or different error detection and correction capability than others. An embodiment is shown which handles both unencoded parity-type detection/correction and more multiple types of more powerful methods, such as Reed-Solomon encoding, in a transparent, real-time fashion, by packing the data words differently for each case. Moreover, the processor loading involved in these error-correction techniques can be spread out in time, so that not all channels need to be handled at the same time. This is accomplished by staggering the beginning times of different data messages.

20           In another embodiment, the at least one service unit includes a service modem for upstream transmission of upstream telephony information and upstream control data within a channel band of the second frequency bandwidth corresponding to one of the channel bands of the first frequency bandwidth in which the service modem receives downstream telephony information and downstream control information. Alternatively, the at least one service unit includes a

25           multi-service modem for upstream transmission of upstream telephony information and upstream control data within a plurality of channel bands of the second frequency bandwidth corresponding to a plurality of the channel bands of the first frequency bandwidth in which the multi-service modem receives downstream telephony information and downstream control information.

In still another embodiment, the plurality of control channels of the first frequency bandwidth and the plurality of control channels of the second frequency bandwidth each include at least one synchronization channel.

In other embodiments, different modulation techniques are utilized for different carriers. For example, different modulation techniques are utilized for different telephony channels. As another example, the aforementioned IOC channels may be modulated as differential binary phase-shift keyed (BPSK) signals, while the payload data channels are modulated as 5-bit quadrature amplitude modulated (QAM32) signals, in order to enhance the use of IOC channels for subband tracking, and for other purposes. The constellation defining the modulated signals can be constructed to achieve a minimal number of bit errors for small errors in amplitude or phase of the received signal; broadly, the constellation points are mapped to bit combinations in a scheme analogous to a Gray code.

A communication system which addresses the problems inherent in the system, in particular, ingress problems is also described. The communication system includes a distribution network between a head end terminal and at least one remote unit. The head end terminal receives upstream telephony information and upstream control data in a frequency bandwidth over the distribution network. The head end terminal includes a head end multicarrier demodulator for demodulating at least upstream telephony information modulated on a plurality of orthogonal carriers in the frequency bandwidth. The demodulator includes at least one polyphase filter for filtering the at least upstream telephony information modulated on the plurality of orthogonal carriers to provide ingress protection for the modulated orthogonal carriers. The head end terminal also includes a head end controller operatively connected to the head end multicarrier demodulator for controlling receipt of the upstream control data and upstream telephony information. The system further includes at least one service unit modulator, each service unit modulator associated with at least one remote unit and operatively connected to the distribution network for modulating at least upstream telephony information on at least one carrier

**S**

10

15

20

25

5

10

5

20

25

In another alternate method embodiment, the method comprises the steps of sampling the parity bit of the n-bit channel and deriving a probable bit error rate

from the sampling of the parity bit over a first time period. The probable bit error rate over the first time period is compared to a first predetermined bit error rate value to determine if the n-bit channel is corrupted. A probable bit error rate from the sampling of the parity bit over a second time period is derived. The second time period is longer than the first time period and runs concurrently therewith. The probable bit error rate over the second time period is compared to a second predetermined bit error rate value to determine if the n-bit channel is corrupted.

In yet another alternate embodiment, a method for monitoring at least one telephony communication channel includes equalizing a signal on the channel and monitoring the equalization of the signal to produce a probable bit error rate as a function of the equalization.

In still yet another alternate embodiment, a method for monitoring at least one unallocated telephony communication channel includes periodically monitoring the at least one unallocated telephony communication channel. Error data for the at least one unallocated telephony communication channel accumulated and the at least one unallocated telephony communication channel is allocated based on the error data.

A multi-point to point communication system utilizing a distributed loop method is also described. The communication system in accordance with the present invention includes a distribution network and a head end terminal for downstream transmission of downstream control data and downstream telephony information in a first frequency bandwidth over the distribution network. The head end terminal receives upstream telephony information and upstream control data in a second frequency bandwidth over the distribution network. The head end terminal further includes a head end multicarrier modem for modulating at least downstream telephony information on a plurality of orthogonal carriers in the first frequency bandwidth and demodulating at least upstream telephony information modulated on a plurality of orthogonal carriers in the second frequency bandwidth. A head end controller is operatively connected to the head end multicarrier modem for

controlling transmission of the downstream telephony information and downstream control data and for controlling receipt of the upstream control data and upstream telephony information. The system includes a plurality of service units. Each service unit is associated with at least one remote unit and operatively connected to the distribution network for upstream transmission of upstream telephony information and upstream control data in the second frequency bandwidth and for receipt of the downstream control data and downstream telephony information in the first frequency bandwidth. Each service unit includes a service unit multicarrier modem for modulating at least the upstream telephony information on at least one carrier orthogonal to at least one other carrier in the second frequency bandwidth and for demodulating at least the downstream telephony information modulated on at least a band of a plurality of orthogonal carriers in the first frequency bandwidth. Each service unit also includes a service unit controller operatively connected to the service unit multicarrier modem for controlling the modulation of and demodulation performed by the service unit multicarrier modem. The service unit controller adjusts at least one local transmission characteristic in response to an adjustment command from the head end controller transmitted in the downstream control data to the at least one remote unit. The head end controller further includes a detector for detecting the at least one local transmission characteristic of the service unit modem associated with the at least one remote unit and for generating the adjustment command as a function of the detected at least one transmission characteristic for transmittal to the service unit associated with the at least one remote unit in the downstream control data.

Distributed system control is also employed for acquiring and tracking remote service units newly connected to (or activated within) the system. An acquisition process makes rough estimates of the frequency, phase, and data-symbol timing of the head-end transmitter, calculates the round-trip delay of data to and from the head end, and then tracks changes in the frequency, phase, and timing of the head end, all with minimal overhead to the transmission of payload data within

the system. A special non-valid data signal is used to signal the start of a training pattern for acquisition purposes. Maintaining accurate power balancing or leveling among the remote units transmitting upstream to the head end is both necessary and difficult in a multipoint-to-point multicarrier system. One embodiment of the invention uses both an automatic gain control or equalizer at the head end and a transmitter output or power control at the remote end to achieve the conflicting goals of wide dynamic range and high resolution amplitude control.

Furthermore, communication system having a distribution network between a head end and a plurality of remote units using a scanning method is described. The system includes the transmission, from the head end, of a plurality of modulated orthogonal carriers having telephony information modulated thereon in a plurality of regions of a first frequency bandwidth. Each of the regions has at least one control channel associated therewith having control information modulated thereon. A scanner at the remote units, scans each of the plurality of regions in the first frequency bandwidth and locks onto the at least one control channel associated with each of the plurality of regions to detect a unique identifier to determine which region of the first frequency bandwidth the remote unit is to tune to and which region in a second frequency bandwidth the remote unit is to transmit within.

In another embodiment, the communication system includes a distribution network between a head end and a plurality of remote units. The head end includes a head end terminal for downstream transmission of downstream control data and downstream telephony information in a first frequency bandwidth over the distribution network and for receipt of upstream telephony information and upstream control data in a second frequency bandwidth over the distribution network. The head end terminal includes a head end multicarrier modem for modulating at least downstream telephony information on a plurality of orthogonal carriers in a plurality of regions of the first frequency bandwidth. The head end multicarrier modem also demodulates at least upstream telephony information modulated on a plurality of orthogonal carriers of a plurality of regions in the second

frequency bandwidth. The plurality of orthogonal carriers in each of the regions includes a plurality of telephony information channels for transmission of telephony information thereon with each of the regions having at least one control channel associated therewith for transmission of control data. The head end terminal also includes a head end controller operatively connected to the head end multicarrier modem for controlling transmission of the downstream telephony information and downstream control data and for controlling receipt of the upstream control data and upstream telephony information. The system further includes a plurality of service unit modems with each service unit modem associated with at least one remote unit and operatively connected to the distribution network for upstream transmission of upstream telephony information and upstream control data in one of the plurality of regions of the second frequency bandwidth and for receipt of the downstream control data and downstream telephony information in one of the plurality of regions in the first frequency bandwidth. Each service unit modem includes a scanner for scanning each of the plurality of regions in the first frequency bandwidth and for locking onto the at least one control channel in each of the plurality of regions to detect a unique identifier for each service unit modem to determine which region of the first frequency bandwidth the service unit modem is to tune to and which region in the second frequency bandwidth the service unit modem is to transmit within.

This invention further provides a three-part RAM structure, the functions of which can be permuted between input, conversion, and output functions. In one embodiment, the conversion RAM section is configured to offer four values to be accessed simultaneously in order to speed operations.

In another embodiment, this invention relates to the field of electronic communication systems, and more specifically to an improved method and apparatus for providing a Sigma-Delta analog-to-digital conversion and decimation for a modem.

According to another aspect of the invention there is provided a data delivery system having a head end connected to a plurality of remote subscribers



over a network, the data originating from general purpose digital computers, wherein the data is transmitted to the remote subscribers in a plurality of data channels, and at the start of a data transmission session, one or more channels are selectively assigned to carry the data between the head end and the subscriber premises with each of the data channels maintaining a substantially constant rate of data transfer between the head end and the destination premise. The system and method further allows that the number of assigned channels assigned to a particular subscriber can be changed from one connection to another to accommodate changes in overall system loading, but at all times maintaining a minimum number of assigned channels so that a minimum rate of data transfer can be maintained between the head end and a subscriber premise. The system also provides asymmetrical operation so that the number of data channels assigned in the downstream path from the head end to the subscribers is much greater in number than the number of upstream data channels.

According to another aspect of the invention there is provided a system of computer data and telephony data transmission over a telecommunications network having a head end connected to a plurality of remote subscribers, the computer data originating from general purpose digital computers, and the telephony data received from or conveyed to the public telephone network. The system comprises transmitting the computer data and telephony data to the remote subscribers in a plurality of data channels, and establishing a computer data or telephony data connection between the head end and a subscriber premise independently of one another. Each of the computer data or telephony data connections are established by assigning one or more of the data channels to carry the computer data and one or more of the data channels to carry the telephony data with at least some of the data channels being available to carry either computer data or telephony data. The system further allows that the number assigned data channels can be changed from one connection to another to the same subscriber so that the overall bandwidth of the network can be reallocated.

The system further allows that the data channels are transmitted using Radio Frequency (RF) signaling, and that the RF is scrambled, thereby providing data security.

According to yet another aspect of the invention, the system allows that a computer data received at the head end can initiate a connection to a remote subscriber, wherein the connection comprises the assignment of one or more data channels to carry the computer data to the subscriber.

According to yet another embodiment, the system 10 or 500 of the present invention includes an ATM modulator which can receive ATM data and modulate it onto the HFC network. In one preferred embodiment, digital video data is delivered over an ATM network, multiplexed and modulated onto the HFC in RF digital OFDM format on data connections established between the head end and a subscriber, as for example described above with respect to system 500. A digital set top box receives the digital video, for example in 4.0 Mbps MPEG or equivalent, and converts it to video for display on a television. A return path over a telephony or data channel allows for interactive digital video.

In another embodiment, a method for transmitting data over a telecommunication system from a head end to a service unit is provided. The service unit is assigned to subband of a transmission channel of the telecommunication system. The subband includes a number of payload channels that transmit data at a first rate and a control channel that transmits data at a second rate. The second rate is slower than the first rate. The system receives a request to transmit data to a service unit at the second, slower rate. The system further determines whether to transmit the data at the first, faster, rate based on the size of the data. When a payload channel in the subband is available to transmit the data at the first rate, the system allocates the payload channel to transmit the data to the service unit over the payload channel at the first rate. When the payload channels are allocated to service units and at least one of the allocated payload channels is idle, the system allocates the idle payload channel to

transmit the data to the service unit over the payload channel at the first rate.

This method can be used to download software and transmission protocols as well as other data that is not time sensitive. Further, the method advantageously provides flexibility in providing bandwidth for transmission of data in the telecommunications system.

In another embodiment, a telecommunications system dynamically allocates bandwidth among a plurality of service units. The system comprises a head end that transmits data over a transmission medium to the service units. The head end includes a modem circuit for narrow band transmission in at least one transmission channel. Each transmission channel includes a number of subbands having a number of payload channels and a control channel in each subband. Further, a control circuit in the head end assigns each service unit to a subband for transmission and receipt of data. The control circuit also allocates a payload channel to a service unit in response to a request for bandwidth for a service unit.

In another embodiment, the control circuit assigns a number of service units to each subband. The control circuit dynamically allocates bandwidth to the service units for selective use of the payload channels in the subband. This increases the number of service units that can be coupled to the system.

In another embodiment, a method for dynamically allocating bandwidth to a service unit in a telecommunications system is provided. The system uses a multi-channel transmission scheme with transmission channels that include a number of subbands. Each subband further includes a number of payload channels. The method begins by receiving a request for a payload channel for a service unit that is assigned to a first subband. The method selects an available payload channel in the first subband and determines if the payload channel is acceptable for providing service to the service unit, e.g., acceptable transmission quality. When the payload channel is acceptable, the method allocates the payload channel to the service unit. When, however, the payload channel is unacceptable, the method selects other channels to find an acceptable payload channel.

In another embodiment, a telecommunications system implements a method for allocating payload channels for a service that use multiple payload channels to communicate with a service unit. The system assigns an identifier for each payload channel that indicates the relative order of the multiple payload channels for the service. The system further monitors the quality of the payload channels of the system. When the quality of one of the multiple payload channels drops below a threshold, the system allocates a different payload channel to replace the original payload channel for the service. Once reallocated, the system uses the identifier for the original payload channel so that the proper order for the allocated payload channels is maintained by the service irrespective of the order that the payload channels are received at the service unit.

In another embodiment, a telecommunications system provides a method for using an upstream payload channel to inform the head end of errors that occur in downstream payload channels. The system monitors a downstream transmission channel at a service unit for transmission errors. Further, the system generates a signal at the service unit that indicates transmission errors in the downstream payload channel. The system also transmits the signal to the head end on an associated upstream payload channel thus allowing the head end to monitor and respond to the performance of the service unit and associated payload channels.

In another embodiment, a method for controlling a plurality of service units in a telecommunication system is provided. The method first assigns an identifier to each service unit. The method further assigns each service unit to a subband of a transmission channel of a narrow band transmission scheme. In the transmission channel, each subband includes a control channel for receiving and transmitting control signals. The method broadcasts the control signals for the service units over the control channels. The method identifies the terminal to use the control signal with the identifier.

In another embodiment, a service unit for use with a communication system transmits signals with a narrow band transmission scheme. The transmission

channel is divided into a number of subbands with each subband including a number of payload channels and a control channel. The service unit includes a modem that is tunable to receive telephony and control signals on a subband of a transmission channel. The service unit further includes a controller circuit coupled to the modem to receive control signals over the control channel and to determine which control signals to use to control the operation of the modem. The service unit also includes interface circuits coupled to the controller for providing signals to a channel unit.

In another embodiment, a method for controlling power usage at a service unit of a telecommunications system is provided. The method comprises determining the type of service supported by each line of a service unit. When the service unit supports analog telephony service, the method determines the hook status of all of the lines of the service unit. When the lines are on-hook, the method powers down the service units to conserve power usage until a request is received to use a line of the service unit.

The present invention describes a method of establishing communication between a head end and a plurality of remote units in a multi-point to point communication system, such as when a fault as described above has left many users of the system without service. The method includes transmitting information from the head end to the plurality of remote units in a plurality of regions of a first frequency bandwidth. Each of the regions has at least one control channel associated therewith. The information transmitted includes identification information corresponding to each of  $n$  remote units of the plurality of remote units. Such information is periodically transmitted for the  $n$  remote units from the head end on the at least one control channel of one of the plurality of regions of the first frequency bandwidth during a first predetermined time period. The identification information for each of the plurality of  $n$  remote units is transmitted out of phase with respect to the identification information for the other of the  $n$  remote units. At each of the  $n$  remote units, the at least one control channel of each of the plurality of regions in the first frequency bandwidth is scanned to detect identification

2025 RELEASE UNDER E.O. 14176

information corresponding to each of the  $n$  remote units to identify a particular region of the plurality of regions that each of the  $n$  remote units is to use for receiving information from the head end.

5 In one embodiment, a region is identified in a second frequency bandwidth in which each of the  $n$  remote units is to transmit within. The method further includes serially performing synchronization for each of the  $n$  remote units for communication with the head end, during a second predetermined time period after the first predetermined time period.

10 A multi-point to point communication system having a distribution network between a head end and a plurality of remote units for accomplishing the above method includes means for transmitting information from the head end to the plurality of remote units in a plurality of regions of a first frequency bandwidth. Each of the regions has at least one control channel associated therewith. The transmitting means further periodically transmits identification information  
15 corresponding to each of a set of  $n$  remote units of the plurality of remote units on at least one control channel of one of the plurality of regions of the first frequency bandwidth during a first predetermined time period of an identification and synchronization time period. The identification information for each of the plurality of  $n$  remote units is transmitted out of phase with respect to the identification information for the other of the  $n$  remote units. The system further includes at each  
20 of the  $n$  remote units, means for scanning the at least one control channel of each of the plurality of regions in the first frequency bandwidth to detect identification information during the first predetermined time period corresponding to each of the  $n$  remote units to identify a particular region of the plurality of regions that each of the  $n$  remote units is to use for receiving information from the head end. Further, at  
25 each of the  $n$  remote units, the system includes means for modulating at least upstream telephony information on at least one carrier in a second frequency bandwidth orthogonal at the head end terminal to at least one other carrier in the second frequency bandwidth and for adjusting at least one local transmission

000034-07004  
000034-07004

characteristic in response to an adjustment command from the head end. Means at the head end for detecting the at least one local transmission characteristic of each of the n remote units and for generating the adjustment commands as a function of the detected at least one transmission characteristic for transmittal to the n remote units to serially perform synchronization for each of the n remote units during a second predetermined time period of the identification and synchronization time period is also included in the system.

The present invention is a hybrid fiber/coax video and telephony communication network which integrates bi-directional telephony and interactive video services into one network including optical fiber and coaxial cable distribution systems. The present invention utilizes optical fiber as the transmission medium for feeding a plurality of optical distribution nodes with video and telephony information from a head end. Coaxial cable distribution systems are utilized for connection of the distribution nodes to a plurality of remote units. The head end optically transmits the video information downstream to the nodes where it is converted to electrical signals for distribution to the remote units. Telephony information is also optically transmitted to the nodes in frequency bandwidths unused by the video information. The downstream telephony and video optical signals are converted to electrical telephony and video signals for distribution to the plurality of remote units. The network provides for transmission of upstream electrical data signals, for example telephony signals, to the head end by transmitting from the remote units upstream electrical data signals to the distribution nodes where such upstream electrical data signals are converted to upstream optical signals for transmission to the head end.

In one embodiment, the head end includes a first distribution terminal having at least one optical transmitter for transmitting optical downstream telephony signals on at least one optical fiber. In addition, the head end includes a second distribution terminal having a separate optical transmitter for transmitting an optical downstream video signal on an optical fiber line.

In another embodiment, the video and telephony signal distribution network transmits optical downstream video and telephony signals on at least one optical fiber in a first frequency bandwidth. In this embodiment, a second frequency bandwidth is reused for transmission of upstream electrical data signals generated at the remote units. The second frequency bandwidth is reused for transmission by each remote unit.

In another embodiment of the invention, a filter is utilized at service units which interface the coaxial distribution systems to user equipment. The ingress filter allows for passage of downstream video signals to video equipment units and blocks downstream telephony signals transmitted in a different frequency bandwidth.

#### Description of the Drawings

- Figure 1 shows a block diagram of a communication system in accordance with the present invention utilizing a hybrid fiber/coax distribution network;
- Figure 2 is an alternate embodiment of the system of Figure 1;
- Figure 3 is a detailed block diagram of a host digital terminal (HDT) with associated transmitters and receivers of the system of Figure 1;
- Figure 4 is a block diagram of the associated transmitters and receivers of Figure 3;
- Figure 5 is a block diagram of an optical distribution node of the system of Figure 1;
- Figures 6, 7 are embodiments of frequency shifters for use in the optical distribution node of Figure 5 and the telephony upstream receiver of Figure 4, respectively;
- Figure 8 is a general block diagram of an integrated service unit (ISU) such as a home integrated service unit (HISU) or a multiple integrated



5

10

15

20

25

Figure 25 is a block diagram of a CXTU upstream transmission architecture

5

Figure 27 is a flow diagram of a acquisition distributed loop routine for use with the system of Figure 1;

10

Figure 30 is an enlarged view of part of the magnitude response of Figure 29;

Figure 32 is a block diagram of a polyphase filter structure of the ingress filter structure and FFT of Figure 31;

**20**

Figure 35 is a block diagram of internal equalizer operation for the receiver architectures of the first transport embodiment;

25

Figure 37 is a block diagram of an MCC modem architecture of the CXMU for the second transport embodiment of the system of Figure 1;

Figure 38 is a block diagram of a subscriber modem architecture of the HISU

for the second transport embodiment of the system of Figure 1;

Figure 39 is a block diagram of a modem of the subscriber modem architecture of Figure 38;

Figure 40 is a block diagram for channel monitoring used in the system of Figure 1;

Figures 41, 42, 43 are flow diagrams for error monitor portions of channel monitor routines of Figure 40;

Figure 44 is an alternate flow diagram for the diagram of Figure 42;

Figure 45 is a flow diagram for a background monitor portion of the channel monitor routines of Figure 40;

Figure 46 is a flow diagram for a backup portion of the channel monitor routines of Figure 40;

Figures 47, 48 are a flow diagram of an acquisition distributed loop routine for use with another embodiment of the system of Figure 1;

Figure 49 is a flow diagram of a downstream tracking loop for use with the embodiment of Figures 47 and 48.

Figure 50 is a flow diagram of an upstream tracking loop for use with the embodiment of Figures 47 and 48.

Figure 51 is a block diagram showing the locking of all clocks within a system.

Figures 52, 53 depict phase diagrams of symbol waveforms in an embodiment of the invention.

Figures 54, 55, 56, 57 describe error rates and message-encoding methods for use in a system according to the invention.

Figure 58 is a block diagram of a scrambler for use in the invention.

Figure 59 is a block diagram of a control circuit for a CXMU of an HDT in a telecommunications system;

Figures 60, 61, 62 are flow charts that illustrate methods for assigning subbands and allocating payload channels in a telecommunications system that uses a multi-carrier communication scheme;

FIG. 39

10

15

20

25

Figure 68 is a flow chart that illustrates error monitoring by the channel manager,

5

Figure 71 is a block diagram of modem 2400 which includes a FFT system 2100 configured to perform an IFFT in transmitter section 2401 and another FFT system 2100 configured to perform an FFT in receiver section 2402;

10

Figure 74 is a block diagram of one embodiment of a dual radix core 2600;

Figures 83, 84, 85, 86, 87, 88, 89, 90 together form a table showing the order of calculations for a "transposed butterfly sub-operation";

25

Figure 93 is a block diagram of modem 2400 which includes a Sigma-Delta

ADC and decimator system to drive FFT system 2100;

- Figure 94 is a more detailed block diagram of modem receiver 2402;
- Figure 95 is a detailed block diagram of one embodiment of a Sigma-Delta converter 2840;
- Figure 96 is an overall schematic diagram of the data delivery transport system according to the present invention;
- Figure 97 is a simplified block diagram of the head-end terminal 12 of the system 500 according to the present invention;
- Figure 98 illustrates a Personal Cable Data Modem (PCDM) 540 and a Data Modem Service Module (DMSM) 550;
- Figure 99 illustrates in greater detail a PCDM 540;
- Figure 100 illustrates a Data Modem Channel Unit (DMCU) 560;
- Figure 101 shows a graph of average bandwidth per user as a function of the number of users for the system 500 according to the present invention;
- Figure 102 is a simplified block diagram of the data transport and framing of the system 500 according to the present invention;
- Figure 103 illustrates a Local Area Network Unit (LANU) 580 according to the present invention;
- Figure 104 illustrates in more detail a DMSM 550 according to the present invention;
- Figure 105 illustrates in more detail a DMCU 560 according to the present invention;
- Figures 106, 107, 108, 109 illustrate the call setup for a data connection on the system 500 according to the present invention;
- Figure 110 illustrates a call termination sequence on the system 500 according to the present invention;
- Figure 111 illustrates the software of a LANU 580 according to the present

invention;

Figure 112 illustrates a PCDM 620 adapted for asymmetrical data delivery;

Figure 113 illustrates the head-end configuration for asymmetrical data delivery according to the present invention;

Figures 114, 115 illustrate another alternate embodiment of the invention wherein digital video is received over an ATM network and transmitted over a modified form of system 10/500;

Figure 116 shows a block diagram of a hybrid fiber/coax network in accordance with the present invention;

Figure 117 is a block diagram of a head end host distribution terminal of the network of Figure 116;

Figure 118 is a block diagram of an optical distribution node of the network of Figure 116;

Figure 119 is a block diagram of a home coaxial line unit of the network of Figure 116;

Figure 120 is a block diagram of an alternative embodiment for transmission from the head end to the optical distribution nodes in accordance with the present invention;

Figure 121 is a block diagram of an impulse shaping technique utilized in accordance with the present invention;

Figure 122 is a block diagram of an alternative embodiment of the optical to electrical converter of the head end host distribution terminal of Figure 117;

Figure 123 is a block diagram of an alternative embodiment of the head end host distribution terminal of Figure 117.

#### Detailed Description of the Preferred Embodiment

The communication system 10, as shown in Figure 1, of the present invention is an access platform primarily designed to deliver residential and

business telecommunication services over a hybrid fiber-coaxial (HFC) distribution network 11. The system 10 is a cost-effective platform for delivery of telephony and video services. Telephony services may include standard telephony, computer data and/or telemetry. In addition, the present system is a flexible platform for accommodating existing and emerging services for residential subscribers.

The hybrid fiber-coaxial distribution network 11 utilizes optical fiber feeder lines to deliver telephony and video service to a distribution node 18 (referred to hereinafter as the optical distribution node (ODN)) remotely located from a central office or a head end 32. From the ODNs 18, service is distributed to subscribers via a coaxial network. Several advantages exist by utilizing the HFC-based communication system 10. By utilizing fiber installed in the feeder, the system 10 spreads the cost of optoelectronics across hundreds of subscribers. Instead of having a separate copper loop which runs from a distribution point to each subscriber ("star" distribution approach), the system 10 implements a bused approach where a distribution coaxial leg 30 passes each home and subscribers "tap" the distribution coaxial leg 30 for service. The system 10 also allows non-video services to be modulated for transmission using more cost-effective RF modem devices in dedicated portions of the RF spectrum. Finally, the system 10 allows video services to be carried on existing coaxial facilities with no additional subscriber equipment because the coaxial distribution links can directly drive existing cable-ready television sets.

It should be apparent to one skilled in the art that the modem transport architecture described herein and the functionality of the architecture and operations surrounding such architecture could be utilized with distribution networks other than hybrid fiber coax networks. For example, the functionality may be performed with respect to wireless systems. Therefore, the present invention contemplates use of such systems in accordance with the accompanying claims.

The system 10 includes host digital terminals 12 (HDTs) which implement all common equipment functions for telephony transport, such as network interface,

synchronization, DS0 grooming, and operations, administration, maintenance and provisioning (OAM&P) interfaces, and which include the interface between the switching network and a transport system which carries information to and from customer interface equipment such as integrated service units 100 (ISUs).

Integrated services units (ISUs) 100, such as home integrated service units (HISUs) 68 or multiple user integrated service units (MISUs) 66, which may include a business integrated service unit as opposed to a multiple dwelling integrated service unit, implement all customer interface functions and interface to the transport system which carries information to and from the switched network. In the present system, the HDT 12 is normally located in a central office and the ISUs 100 are remotely located in the field and distributed in various locations. The HDT 12 and ISUs 100 are connected via the hybrid fiber-coax distribution network 11 in a multi-point to point configuration. In the present system, the modem functionality required to transport information over the HFC distribution network 11 is performed by interface equipment in both the HDT 12 and the ISUs 100. Such modem functionality is performed utilizing orthogonal frequency division multiplexing.

The communication system shall now be generally described with reference to Figures 1, 3 and 8. The primary components of system 10 are host digital terminals (HDTs) 12, video host distribution terminal (VHDT) 34, telephony downstream transmitter 14, telephony upstream receiver 16, the hybrid fiber coax (HFC) distribution network 11 including optical distribution node 18, and integrated service units 66, 68 (shown generally as ISU 100 in Figure 8) associated with remote units 46. The HDT 12 provides telephony interface between the switching network (noted generally by trunk line 20) and the modem interface to the HFC distribution network for transport of telephony information. The telephony downstream transmitter 14 performs electrical to optical conversion of coaxial RF downstream telephony information outputs 22 of an HDT 12, shown in Figure 3, and transmits onto redundant downstream optical feeder lines 24. The telephony upstream receiver 16 performs optical to electrical conversion of optical signals on



redundant upstream optical feeder lines 26 and applies electrical signals on coaxial RF upstream telephony information inputs 28 of HDT 12. The optical distribution node (ODN) 18 provides interface between the optical feeder lines 24 and 26 and coaxial distribution legs 30. The ODN 18 combines downstream video and telephony onto coaxial distribution legs 30. The integrated services units provide modem interface to the coaxial distribution network and service interface to customers.

The HDT 12 and ISUs 100 implement the telephony transport system modulator-demodulator (modem) functionality. The HDT 12 includes at least one RF MCC modem 82, shown in Figure 3 and each ISU 100 includes an RF ISU modem 101, shown in Figure 8. The MCC modems 82 and ISU modems 101 use a multi-carrier RF transmission technique to transport telephony information, such as DS0+ channels, between the HDT 12 and ISUs 100. This multi-carrier technique is based on orthogonal frequency division multiplexing (OFDM) where a bandwidth of the system is divided up into multiple carriers, each of which may represent an information channel. Multi-carrier modulation can be viewed as a technique which takes time-division multiplexed information data and transforms it to frequency-division multiplexed data. The generation and modulation of data on multiple carriers is accomplished digitally, using an orthogonal transformation on each data channel. The receiver performs the inverse transformation on segments of the sampled waveform to demodulate the data. The multiple carriers overlap spectrally. However, as a consequence of the orthogonality of the transformation, the data in each carrier can be demodulated with negligible interference from the other carriers, thus reducing interference between data signals transported. Multi-carrier transmission obtains efficient utilization of the transmission bandwidth, particularly necessary in the upstream communication of a multi-point to point system. Multi-carrier modulation also provides an efficient means to access multiple multiplexed data streams and allows any portion of the band to be accessed to extract such multiplexed information, provides superior noise immunity to impulse noise as a

consequence of having relatively long symbol times, and also provides an effective means for eliminating narrowband interference by identifying carriers which are degraded and inhibiting the use of these carriers for data transmission (such channel monitoring and protection is described in detail below). Essentially, the telephony transport system can disable use of carriers which have interference and poor performance and only use carriers which meet transmission quality targets.

Further, the ODNs 18 combine downstream video with the telephony information for transmission onto coaxial distribution legs 30. The video information from existing video services, generally shown by trunk line 20, is received by and processed by head end 32. Head end 32 or the central office, includes a video host distribution terminal 34 (VHDT) for video data interface. The VHDT 34 has optical transmitters associated therewith for communicating the video information to the remote units 46 via the ODNs 18 of the HFC distribution network 11.

The telephony transmitter 14 of the HDTs 12, shown in Figure 3 and 4, includes two transmitters for downstream telephony transmission to protect the telephony data transmitted. These transmitters are conventional and relatively inexpensive narrow band laser transmitters. One transmitter is in standby if the other is functioning properly. Upon detection of a fault in the operating transmitter, the transmission is switched to the standby transmitter. In contrast, the transmitter of the VHDT 34 is relatively expensive as compared to the transmitters of HDT 12 as it is a broad band analog DFB laser transmitter. Therefore, protection of the video information, a non-essential service unlike telephony data, is left unprotected. By splitting the telephony data transmission from the video data transmission, protection for the telephony data alone can be achieved. If the video data information and the telephony data were transmitted over one optical fiber line by an expensive broad band analog laser, economies may dictate that protection for telephony services may not be possible. Therefore, separation of such transmission is of importance.

Further with reference to Figure 1, the video information is optically transmitted downstream via optical fiber line 40 to splitter 38 which splits the optical video signals for transmission on a plurality of optical fiber lines 42 to a plurality of optical distribution nodes 18. The telephony transmitter 14 associated with the HDT 12 transmits optical telephony signals via optical fiber feeder line 42 to the optical distribution nodes 18. The optical distribution nodes 18 convert the optical video signals and optical telephony signals for transmission as electrical outputs via the coaxial distribution portion of the hybrid fiber coax (HFC) distribution network 11 to a plurality of remote units 46. The electrical downstream video and telephony signals are distributed to ISUs via a plurality of coaxial legs 30 and coaxial taps 44 of the coaxial distribution portion of the HFC distribution network 11.

The remote units 46 have associated therewith an ISU 100, shown generally in Figure 8, that includes means for transmitting upstream electrical data signals including telephony information, such as from telephones and data terminals, and in addition may include means for transmitting set top box information from set top boxes 45 as described further below. The upstream electrical data signals are provided by a plurality of ISUs 100 to an optical distribution node 18 connected thereto via the coaxial portion of the HFC distribution network 11. The optical distribution node 18 converts the upstream electrical data signals to an upstream optical data signal for transmission over an optical fiber feeder line 26 to the head end 32.

Figure 2 generally shows an alternate embodiment for providing transmission of optical video and optical telephony signals to the optical distribution nodes 18 from head end 32, the HDT 12 and VHDT 34 in this embodiment utilize the same optical transmitter and the same optical fiber feeder line 36. The signals from HDT 12 and VHDT 34 are combined and transmitted optically from headend 32 to splitter 38. The combined signal is then split by splitter 38 and four split signals are provided to the optical distribution nodes 18 for distribution to the

remote units by the coaxial distribution legs 30 and coaxial taps 44. Return optical telephony signals from the ODNs 18 would be combined at splitter 38 for provision to the headend. However, as described above, the optical transmitter utilized would be relatively expensive due to its broad band capabilities, lessening the probabilities of being able to afford protection for essential telephony services.

As one skilled in the art will recognize, the fiber feeder lines 24, 26, as shown in Figure 1, may include four fibers, two for transmission downstream from downstream telephony transmitter 14 and two for transmission upstream to upstream telephony receiver 16. With the use of directional couplers, the number of such fibers may be cut in half. In addition, the number of protection transmitters and fibers utilized may vary as known to one skilled in the art and any listed number is not limiting to the present invention as described in the accompanying claims.

The present invention shall now be described in further detail. The first part of the description shall primarily deal with video transport. The remainder of the description shall primarily be with regard to telephony transport.

#### VIDEO TRANSPORT

The communication system 10 includes the head end 32 which receives video and telephony information from video and telephony service providers via trunk line 20. Head end 32 includes a plurality of HDTs 12 and a VHDT 34. The HDT 12 includes a network interface for communicating telephony information, such as T1, ISDN, or other data services information, to and from telephony service providers, such communication also shown generally by trunk line 20. The VHDT 34 includes a video network interface for communicating video information, such as cable TV video information and interactive data of subscribers to and from video service providers, such communication also shown generally by trunk line 20.

The VHDT 34 transmits downstream optical signals to a splitter 38 via video optical fiber feeder line 40. The passive optical splitter 38 effectively makes four copies of the downstream high bandwidth optical video signals. The duplicated

downstream optical video signals are distributed to the correspondingly connected optical distribution nodes 18. One skilled in the art will readily recognize that although four copies of the downstream video signals are created, any number of copies may be made by an appropriate splitter and that the present invention is not limited to any specific number.

The splitter is a passive means for splitting broad band optical signals without the need to employ expensive broad band optical to electrical conversion hardware. Optical signal splitters are commonly known to one skilled in the art and available from numerous fiber optic component manufacturers such as Gould, Inc. In the alternative, active splitters may also be utilized. In addition, a cascaded chain of passive or active splitters would further multiply the number of duplicated optical signals for application to an additional number of optical distribution nodes and therefore increase further the remote units serviceable by a single head end. Such alternatives are contemplated in accordance with the present invention as described by the accompanying claims.

The VHDT 34 can be located in a central office, cable TV head end, or a remote site and broadcast up to about 112 NTSC channels. The VHDT 34 includes a transmission system like that of a LiteAMp™ system available from American Lightwave Systems, Inc., currently a subsidiary of the assignee hereof. Video signals are transmitted optically by amplitude modulation of a 1300 nanometer laser source at the same frequency at which the signals are received (i.e. the optical transmission is a terahertz optical carrier which is modulated with the RF video signals). The downstream video transmission bandwidth is about 54-725 MHz. One advantage in using the same frequency for optical transmission of the video signal as the frequency of the video signals when received is to provide high bandwidth transmission with reduced conversion expense. This same-frequency transmission approach means that the modulation downstream requires optical to electrical conversion or proportional conversion with a photodiode and perhaps amplification, but no frequency conversion. In addition, there is no sample data

bandwidth reduction and little loss of resolution.

An optical distribution node 18, shown in further detail in Figure 5, receives the split downstream optical video signal from the splitter 38 on optical fiber feeder line 42. The downstream optical video signal is applied to a downstream video receiver 400 of the optical distribution node 18. The optical video receiver 400 utilized is like that available in the Lite AMp™ product line available from American Lightwave Systems, Inc. The converted signal from video receiver 400, proportionally converted utilizing photodiodes, is applied to bridger amplifier 403 along with converted telephony signals from downstream telephony receiver 402. The bridger amplifier 403 simultaneously applies four downstream electrical telephony and video signals to duplex filters 406 which allow for full duplex operation by separating the transmit and receive functions when signals of two different frequency bandwidths are utilized for upstream and downstream transmission. There is no frequency conversion performed at the ODN 18 with respect to the video or the downstream telephony signals as the signals are passed through the ODNs to the remote units via the coaxial portion of the HFC distribution network 11 in the same frequency bandwidth as they are received at the ODNs 18.

After the ODN 18 has received the downstream optical video signals and such signals are converted to downstream electrical video signals, the four outputs of the ODN 18 are applied to four coaxial legs 30 of the coaxial portion of the HFC distribution network 11 for transmission of the downstream electrical video signals to the remote units 46. Such transmission for the electrical video signals occurs in about the 54-725 MHz bandwidth. Each ODN 18 provides for the transmission on a plurality of coaxial legs 30 and any number of outputs is contemplated in accordance with the present invention as described in the accompanying claims.

As shown in Figure 1, each coaxial cable leg 30 can provide a significant number of remote units 46 with downstream electrical video and telephony signals through a plurality of coaxial taps 44. Coaxial taps are commonly known to one

skilled in the art and act as passive bidirectional pickoffs of electrical signals. Each coaxial cable leg 30 may have a number of coaxial taps 44 connected in series. In addition, the coaxial portion of the HFC distribution network 11 may use any number of amplifiers to extend the distance data can be sent over the coaxial portion of such HFC distribution network 11.

Downstream video signals are provided from the coaxial taps 44 to the remote units 46. The video signal from the coaxial tap 44 is provided to an HISU 68 which is generally shown by the block diagram of ISU 100 in Figure 8. The ISU 100 is provided with the downstream electrical video and telephony signal from tap 44 and it is applied to duplex filter 104. The downstream electrical video and telephony signal is passed through the duplex filter 104 to both an ingress filter 105 and ISU modem 101. The downstream video signal is passed by the ingress filter 105 to video equipment via an optional set top box 45. The downstream electrical telephony signal applied from the duplex filter 104 to the ISU modem 101 is processed as described in further detail below.

Ingress filter 105 provides the remote unit 46 with protection against interference of signals applied to the video equipment as opposed to those provided to other user equipment such as telephones or computer terminals. Ingress filter 105 passes the video signals; however, it blocks those frequencies not utilized by the video equipment. By blocking those frequencies not used by the video equipment, stray signals are eliminated that may interfere with the other services by the network to at least the same remote unit.

The set top box 45 is an optional element at the remote unit 46. Interactive video data from set top box 45 would be transmitted by an additional separate RF modem provided by the video service provider at a relatively low frequency in the bandwidth of about 5 to 40 MHz. Such frequency must not be one used for the transport of upstream and downstream telephony data and downstream video.

For an MISU 66, a separate coaxial line from coaxial tap 44 is utilized to provide transmission of video signals from the coaxial tap 44 to the set top box 45

and thus for providing downstream video signals to video equipment 47. The ingress filter 105 as shown in Figure 8 is not a part of the MISU 66 as indicated by its dashed representation.

Alternative embodiments of the VHDT 34 may employ other modulation and mixing schemes or techniques to shift the video signals in frequency, and other encoding methods to transmit the information in a coded format. Such techniques and schemes for transmitting analog video data, in addition to those transmitting digital video data, are known to one skilled in the art and are contemplated in accordance with the spirit and scope of the present invention as described in the accompanying claims.

#### TELEPHONY TRANSPORT

With reference to Figure 3, telephony information and ISU operations and control data (hereinafter referred to as control data) modulated on carriers by MCC modem 82 is transmitted between the HDT 12 and the telephony downstream transmitter 14 via coaxial lines 22. Telephony information and control data modulated on carriers by ISUs 100 is received at telephony upstream receiver 16 and communicated to the MCC modem 82 via coaxial cable lines 28. The telephony downstream transmitter 14 and the telephony upstream receiver 16 transmit and receive, respectively, telephony information and control data via optical fiber feeder lines 24 and 26 to and from a corresponding optical distribution node 18. The control data may include all operations, administration, maintenance & provisioning (OAM&P) for providing the telephony services of the system 10 and any other control data necessary for providing transport of telephony information between the HDT 12 and the ISUs 100.

A block diagram of the HDT 12 is shown in Figure 3. The HDT 12 includes the following modules: Eight DS1 Units (DS1U) (seven quad-DS1 units 48 plus one protection unit 50), one protection switch & test conversion unit 52 (PSTU), two clock & time slot interchange units 54 (CTSU) (one active and one



standby/protection unit), six coax master units 56 (CXMUs) (three active and three standby/protection units), two shelf control units 58 (SCNUs) (one active and one standby/protection unit), and two power supply units 60 (PWRUs) (two load-sharing units which provide the appropriate HDT voltages from a central office supply). The DS1U units can also be adapted to transfer data in the standard E1U format, if desired.

The HDT 12 comprises all the common equipment functions of the telephony transport of the communication system 10. The HDT 12 is normally located in a central office and directly interfaces to a local digital switch or digital network element equipment. The HDT provides the network interface 62 for all telephony information. Each HDT accommodates from 2 to 28 DSX-1 inputs at the network interface 62, representing a maximum of 672 DS0 channels. The HDT 12 also provides all synchronization for telephony transport in the system 10. The HDT 12 may operate in any one of three synchronization modes: external timing, line timing or internal timing. External timing refers to synchronization to a building integrated timing supply reference which is sourced from a central office in which the HDT 12 is located. Line timing is synchronized to the recovered clock from a DSX-1 signal normally derived from the local digital switch. Internal timing is a free-running or hold-over operation where the HDT maintains its own synchronization in the absence of any valid reference inputs.

The HDT 12 also provides quarter-DS0 grooming capabilities and implements a 4096 x 4096 full-access, non-blocking quarter-DS0 (16 kbps) cross-connect capability. This allows DS0s and quarter-DS0s (ISDN "D" channels) to be routed from any timeslot at the DSX-1 network interface 62 to any customer serviced by any ISU 100.

The HDT 12 further provides the RF modem functionality required for telephony transport over the HFC distribution network 11 including the MCC modem 82. The HDT 12 accommodates up to three active CXMUs 56 for providing the modem interface to the HFC distribution network 11 and also provides one-for-

5

10

## 15

20

25

preferably within a distance of 100 feet or less. The downstream telephony transmitter 14 receives the coaxial RF outputs from the active CXMUs 56, each within a 6 MHz frequency band, and combines them at combiner 25 into a single RF signal. Each 6 MHz frequency band is separated by a guard band as is known to one skilled in the art. Downstream telephony information is then transmitted in about the 725-800 MHz frequency band. The telephony transmitter 14 passes the combined signal through a 1-to-2 splitter (not shown), thereby producing redundant downstream electrical signals. The two redundant signals are each delivered to redundant laser transmitters 501 for electrical-to-optical conversion and the redundant signals modulate an optical output such that the output of the downstream telephony transmitter 14 is on two optical feeder lines 24, each having an identical signal modulated thereon. This provides protection for the downstream telephony portion of the present system. Both Fabry-Perot lasers in the telephony transmitter 14 are active at all times. All protection functions are provided at the receive end of the optical transmission (located at the ODN 18) where one of two receivers is selected as "active;" therefore, the telephony transmitter 14 requires no protection switching capabilities.

#### Upstream Telephony Receiver

The upstream telephony receiver 16 performs the optical-to-electrical conversion on the upstream optical telephony signals on the upstream optical feeder lines 26 from the ODN 18. The upstream telephony receiver 16 is normally co-located in the central office with the HDT 12, and provides an electrical coaxial output to the HDT 12, and a coaxial output 23 to be provided to a video set-top controller (not shown). Upstream telephony information is routed via coax lines 28 from the upstream telephony receiver 16 to active CXMUs 56 of the HDT 12. The coaxial link 28 between the HDT 12 and the upstream telephony receiver 16 is preferably limited to a distance of 100 feet or less and is an intra-office link. Video set-top controller information, as described in the Video Transport section hereof, is

located in a bandwidth of the RF spectrum of 5-40 MHz which is not utilized for upstream telephony transport such that it is transmitted along with the upstream telephony information.

The upstream telephony receiver 16 has dual receivers 502 for the dual upstream optical fiber feeders lines 26. These feeder lines 26 carry redundant signals from the ODN 18 which contain both telephony information and control data and also video set-top box information. The upstream telephony receiver 16 performs automatic protection switching on the upstream feeder lines 26 from the ODN. The receiver 502 selected as "active" by protection logic is split to feed the coaxial outputs 28 which drive the HDT 12 and output 23 is provided to the set-top controller (not shown).

#### Optical Distribution Node

Referring to Figure 5, the ODN 18 provides the interface between the optical feeder lines 24 and 26 from the HDT 12 and the coaxial portion of the HFC distribution network 11 to the remote units 46. As such, the ODN 18 is essentially an optical-to-electrical and electrical-to-optical converter. The maximum distance over coax of any ISU 100 from an ODN 18 is preferably about 6 km and the maximum length of the combined optical feeder line/coaxial drop is preferably about 20 km. The optical feeder line side of the ODN 18 terminates six fibers although such number may vary. They include: a downstream video feeder line 42 (single fiber from video splitter 38), a downstream telephony feeder line 24 (from downstream telephony transmitter 14), a downstream telephony protection feeder line 24 (from downstream telephony transmitter 14), an upstream telephony feeder line 26 (to upstream telephony receiver 16), an upstream protection feeder line 26 (to upstream telephony receiver 16), and a spare fiber (not shown). The ODN 18 provides protection switching functionality on the receive optical feeder lines 24 from the downstream telephony transmitter. The ODN provides redundant transmission on the upstream optical feeder lines 26 to the upstream telephony

receiver. Protection on the upstream optical feeder lines is controlled at the upstream telephony receiver 16. On the coaxial distribution side of ODN 18, the ODN 18 terminates up to four coaxial legs 30.

In the downstream direction, the ODN 18 includes downstream telephony receiver 402 for converting the optical downstream telephony signal into an electrical signal and a bridger amplifier 403 that combines it with the converted downstream video signal from downstream video receiver 400 terminated at the ODN 18 from the VHDT 34. This combined wide-band electrical telephony/video signal is then transported in the spectrum allocated for downstream transmission, for example, the 725-800 MHz band, on each of the four coaxial legs of the coaxial portion of the HFC distribution network 11. As such, this electrical telephony and video signal is carried over the coaxial legs 30 to the ISUs 100; the bridger amplifier 403 simultaneously applying four downstream electrical telephony and video signals to diplex filters 406. The diplex filters 406 allow for full duplex operation by separating the transmit and receive functions when signals at two different frequency bandwidths are utilized for upstream and downstream transmission. There is no frequency conversion available at the ODN 18 for downstream transport as the telephony and video signals are passed through the ODN 18 to the remote units 46 via the coaxial portion of HFC distribution network 11 in the same frequency bandwidth as they are received at the ODN 18. As shown in Figure 1, each coaxial leg 30 can provide a significant number of remote units 46 with downstream electrical video and telephony signals through a plurality of coaxial taps 44. Coaxial taps 44 commonly known to one skilled in the art act as passive bidirectional pickoffs of electrical signals. Each coaxial leg 30 may have a number of coaxial taps connected in a series. In addition, the coaxial portion of the HFC distribution network 11 may use any number of amplifiers to extend the distance data can be sent over the coaxial portions of the system 10. The downstream electrical video and telephony signals are then provided to an ISU 100 (Figure 8), which, more specifically, may be an HISU 68 or an MISU 66 as shown in Figure 1.

5

25

In operation, frequency shifter 64' shifts a block of RF signals from a first frequency range to a second frequency range. For example, as mentioned above, the RF signals provided to frequency shifter may comprise RF signals in the range from

5 to 40 MHz. In one embodiment, ODN 18 comprises three frequency shifters 64'. In this embodiment, the local oscillators 702 of the three frequency shifters provide signals of 76 MHz, 149 MHz, and 222 MHz, respectively. Thus, frequency shifters 64' respectively shift the upstream RF signals approximately to the 50 to 100 MHz, 125 to 175 MHz and 200 to 250 MHz ranges.

If the upstream telephony and set top box signals are upshifted at the ODN 18, the upstream telephony receiver 16 includes frequency shifters 31' to downshift the signals according to the upshifting done at the ODN 18. A combiner 33 then combines the downshifted signals for application of a combined signal to the HDT 12. Such downshifting and combining is only utilized if the signals are upshifted at the ODN 18.

Figure 7 illustrates an embodiment of a frequency shifter, indicated generally at 31', for use in telephony upstream receiver 16 of Figure 8. Frequency shifter 31' returns a block of RF signals shifted by frequency shifter 64' to original frequency range of the block. For example, frequency shifter 31' may return a block of RF signals to 5 to 40 MHz from 50 to 100 MHz.

As discussed in more detail below, the upstream telephony signals processed by frequency shifters 31' and 64' are typically OFDM signals. Thus, frequency shifters 64' must return the RF signals to the original frequency range without introducing adverse phase and frequency errors. To reduce the likelihood of this corruption of the OFDM signals, frequency shifter 31' locks its local oscillator to the local oscillator of a corresponding frequency shifter 64' using a pilot tone transmitted from ODN 18 to telephony upstream receiver 16.

Frequency shifter 31' includes a bandpass filter 706 that is coupled to receive an RF signal from ODN 18. Bandpass filter 706 is coupled to a splitter 708. Splitter 708 is coupled to provide the RF signal to an input of mixer 718. Further, splitter 708 provides a second output that is used to generate a local oscillator signal for mixer 718. This local oscillator signal is phase locked with a corresponding local oscillator 702 of frequency converter 64'. This second output of splitter 708 is

coupled to phase detector 712 through bandpass filter 710. Phase detector 712 is coupled to provide a control signal to voltage controlled oscillator 714. Voltage controlled oscillator 714 is coupled through splitter 716 to provide the local oscillator signal to mixer 718. Splitter 716 further provides a feedback signal to phase detector 712.

In operation, phase detector 712 phase locks local oscillator signal of frequency shifter 31' with local oscillator 702 of a corresponding frequency shifter 64'. Phase detector 712 compares the pilot tone from ODN 18 with the feedback signal from voltage controlled oscillator 714 to generate the control signal for voltage controlled oscillator 714. Consequently, the local oscillator signal provided to mixer 718 is phase locked with the corresponding local oscillator 702 of frequency shifter 64'. Mixer 718 uses the local oscillator signal from splitter 716 and voltage controlled oscillator 714 to shift the block of RF signals received by frequency shifter 31' to the original frequency range of the block of RF signals. Advantageously, unacceptable modifications of the OFDM upstream signal by frequency shifters 64' and 31' are thus avoided.

#### Integrated Services Unit (ISUs)

Referring to Figure 1, the ISUs 100, such as HISU 68 and MISU 66, provide the interface between the HFC distribution network 11 and the customer services for remote units 46. Two basic types of ISUs are shown, which provide service to specific customers. Multiple user integrated service unit 66 (MISUs) may be a multiple dwelling integrated service unit or a business integrated service unit. The multiple dwelling integrated service unit may be used for mixed residential and business environments, such as multi-tenant buildings, small businesses and clusters of homes. These customers require services such as plain old telephone service (POTS), data services, DS1 services, and standard TR-57 services. Business integrated service units are designed to service business environments. They may require more services, for example, data services, ISDN, DS1 services, higher band-



5

10

25

5

10

25

necessary to detect customer line seizure and line idle conditions and transmit these indications to the HDT 12. ISUs 100 terminate and receive control data from the HDT 12 and process the control data received therefrom. Included in this processing are messages to coordinate dynamic channel allocation in the communication system 10. Finally, ISUs 100 generate ISU operating voltages from a power signal received over the HFC distribution network 11 as shown by the power signal 109 taken from diplex filter/tap 104.

#### Data Path in HDT

The following is a detailed discussion of the data path in the host digital terminal (HDT) 12. Referring to Figure 3, the data path between the network facility at the network interface 62 and the downstream telephony transmitter 14 proceeds through the DS1U 48, CTSU 54, and CXMU 56 modules of the HDT 12, respectively, in the downstream direction. Each DS1U 48 in the HDT 12 takes four DS1s from the network and formats this information into four 24-channel, 2.56 Mbps data streams of modified DS0 signals referred to as CTSU inputs 76. Each DS0 in the CTSU input has been modified by appending a ninth bit which can carry multiframe timing, signaling information and control/status messages (Figure 9). This modified DS0 is referred to as a "DS0+." The ninth bit signal (NBS) carries a pattern which is updated each frame and repeats every 24 frames. This maps each 64 kbps DS0 from the network into a 72 kbps DS0+. Thus, the twenty-four DS0 channels available on each DS1 are formatted along with overhead information into twenty-four DS0+ channels on each of four CTSU input streams.

The ninth bit signaling (NBS) is a mechanism developed to carry the multiframe timing, out-of-band signaling bits and miscellaneous status and control information associated with each DS0 between the DS1U and the channel units. Its main functions are to carry the signaling bits to channel units 103 and to provide a multiframe clock to the channel units 103 so that they can insert upstream bit signaling into the DS0 in the correct frame of the multiframe. Because downstream

DS0s may be coming from DS1s which do not share the same multiframe phase each DS0 must carry a multiframe clock or marker which indicates the signaling frames associated with the origination DS1. The NBS provides this capability. Ninth bit signaling is transparent to the OFDM modem transport of the communication system 10.

Up to eight DS1Us 48 may be equipped in a single HDT 12; including seven active DS1Us 48 and a protection DS1U module 50. Thus, 32 CTSU inputs are connected between the DS1Us and the CTSUs 54 but a maximum of 28 can be enabled to carry traffic at any one time. The four remaining CTSU inputs are from either the protection DS1U or a failed DS1U. The PSTU includes switch control for switching the protection DS1U 50 for a failed DS1U.

Each CTSU input is capable of carrying up to 32, 10-bit channels, the first 24 channels carry DS0+s and the remaining bandwidth is unused. Each CTSU input 76 is clocked at 2.56 Mbps and is synchronized to the 8 kHz internal frame signal (Figure 11). This corresponds to 320 bits per 125  $\mu$ sec frame period. These 320 bits are framed as shown in Figure 9. The fourteen gap bits 72 at the beginning of the frame carry only a single activity pulse in the 2nd bit position, the remaining 13 bits are not used. Of the following 288 bits, the first 216 bits normally carry twenty-four DS0+ channels where each DS0+ corresponds to a standard 64 kbps DS0 channel plus the additional 8 kbps signaling bit. Thus, each DS0+ has a bandwidth of 72 kbps (nine bits every 8 kHz frame). The remaining 72 bits are reserved for additional DS0+ payload channels. The final eighteen bits 74 of the frame are unused gap bits.

The clock and time slot interchange unit 54 (CTSU) of the HDT 12 takes information from up to 28 active CTSU input data streams 76 and cross-connects them to up to twenty-four 32-channel, 2.56 Mbps output data streams 78 which are input to the coax master units (CXMUs) 56 of the HDT 12. The format of the data streams between the CTSU 54 and the CXMUs 56 is referred to as a CTSU output. Each CTSU output can also carry up to 32, 10-bit channels like the CTSU input.

The first 28 carry traffic and the remaining bandwidth is unused. Each CTSU output is clocked at 2.56 Mbps and is synchronized to the 8 kHz internal framing signal of the HDT 12 (Figure 11). This corresponds to 320 bits per 125  $\mu$ sec frame period. The frame structure for the 320 bits are as described above for the CTSU input structure.

The HDT 12 has the capability of time and space manipulation of quarter-DS0 packets (16 kbps). This function is implemented with the time slot interchange logic that is part of CTSU 54. The CTSU implements a 4096 x 4096 quarter-DS0 cross-connect function, although not all time slots are utilized. In normal operation, the CTSU 54 combines and relocates up to 672 downstream DS0+ packets (or up to 2688 quarter-DS0 packets) arranged as 28 CTSU inputs of 24 DS0+s each, into 720 DS0+ packets (or 2880 quarter-DS0 packets) arranged as 24 CTSU outputs of 32 DS0+s each.

The system has a maximum throughput of 672 DS0+ packets at the network interface so not all of the CTSU output bandwidth is usable. If more than the 672 channels are assigned on the "CTSU output" side of the CTSU, this implies concentration is being utilized. Concentration is discussed further below.

Each CXMU 56 is connected to receive eight active CTSU outputs 78 from the active CTSU 54. The eight CTSU outputs are clocked by a 2.56 MHz system clock and each carries up to 32 DS0+s as described above. The DS0+s are further processed by the CXMU 56 and a tenth parity bit is appended to each DS0+ resulting in a 10 bit DS0+. These 10 bit packets contain the DS0, the NBS (ninth bit signal) and the parity or data integrity bit (Figure 10). The 10 bit packets are the data transmitted on the HFC distribution network 11 to the ISUs 100. The 10th bit or data integrity bit inserted in the downstream channels is decoded and checked at the ISU and utilized to calculate and generate a parity bit for corresponding channels in the upstream as described above. This upstream parity bit which may be representative of an error in the downstream or upstream channel is utilized to

provide channel protection or monitoring as further described herein.

CTSU 54 is also capable of applying a conventional Reed-Soloman code to transmitted payload data for running error correction. Such codes carry an overhead which must be borne by the inclusion of error-correcting symbols in each block or message transmitted. Generally, two check symbols are required to correct one corrupted data symbol in a message. (The incorrect symbol may contain any number of errors in its 5 bits, as long as all bit errors are confined to the same symbol. But even a single incorrect bit in each of two symbols counts as two errors.) Short messages impose less computational overhead on a system, but can correct fewer errors in the message. Conversely, long messages require more computation and more latency before the corrections can be applied, but their error-correction ability is greater. Figure 54 represents, for an example system, the probability of an uncorrectable error in a frame for various error probabilities in one individual symbol. The solid curve shows the error performance for a 21-frame message having 19 frames of data symbols and two frames of error-correction code; the dashed curve represents a 41-frame message having 37 data and four code frames; the dotted curve gives the best performance, with 73 data frames and eight code frames in an 81-frame message.

The present system allows a choice of different error-correction abilities for different types of data. For example, voice data is highly redundant, and needs little defense against errors. Financial transaction data, on the other hand, wants a large degree of data integrity. In addition, it may be desirable to allow a user to select -- and pay for -- whatever degree of error correction that he desires. CTSU 54, Figure 3, includes a conventional "provisioning table", which specifies a number of parameters relating to particular payload channels. Figure 55 shows a provisioning table 4411 having an added column containing indications for several different amounts of error protection. In method 4410, step 4412 reads the entry for a particular channel to be set up. In this implementation, the entry may specify message lengths of 21, 41, or 81 bits, respectively having the ability to correct 1, 2,

or 4 symbols; the entry may also specify no correction, in which case message blocks do not apply. Step 4413 encodes the table entry in an IOC message and sends it to the ISU whose address appears in that row of table 4111. A general-purpose processor in CXSU 102 of the ISU stores the frame length in step 4414. As the CXSU receives data from modem 101, Figure 8, it decodes the frames of an entire message, 4415, then decodes the check symbols for the message, 4416, and signals an error, 4417, if one exists in the message. Steps 4415-4417 repeat for subsequent messages. The ISU employs the same process to send frames upstream to the head end, using the frame length setting specified in step 4414.

Within both CXSU 102 at the ISU and the CXMU, Figure 12 at the HDT, a 21-frame message or block requires 19 symbol or frame times to decode the message, then has two frames of latency while its two check symbols are decoded. A 41-frame message uses four frames of time for computation of any errors from the four check symbols following its 37 data symbols. An 81-frame message presents any error indication 8 frame times after the end of its 73 data frames. (One extra frame of delay is imposed in the downstream direction due to remapping at the HDT.) If all messages were to start at the same time for all channels in an entire band, the computational load in the HDT would peak during the check-symbol frames, and would be lower at other times. Since the processor must be capable of handling the peak loads, its power is underutilized at other times.

The present system allows a lower-capacity processor to handle error correction by staggering the beginning times of different messages in different subbands of channels, so that not all of them come due at the same time. That is, the start of a message in any channel of a subband is offset from the start of a message or "multiframe" signal, to be described, by a predetermined number of frame times of 125 $\mu$ sec. each. The table below shows how the 24 subbands of Figure 16 are offset, for each message length which can be selected.

Table 1

Subband	21 frames	41 frames	81 frames
0	0	0	0
1	0	0	1
2	0	1	2
3	0	1	3
4	1	2	4
5	1	2	5
6	1	3	6
7	1	3	7
8	2	4	8
9	2	4	9
10	2	5	10
11	2	5	11
12	3	6	12

Only 13 subband settings are required, since no unit tunes more than 130 channels. Giving all 10 channels of each subband the same offset does not overload the processors of the remote units. However, the head end (HDT), which receives and transmits all channels, can enjoy significant relief from not having to encode or decode the check symbols for all channels at the same time.

Figure 56 shows steps 4120 for performing frame staggering. Step 4421 repeats method 4420 for all active payload channels. Step 4422 accesses the current messages for the channels in one subband. Step 4423 calculates the 1, 2, or 4 Reed-Soloman check words for the 21, 41, or 81 message data words. Step 4424 waits N frames past the start of a multiframe, whereupon step 4425 sends the message to modem 82, Figure 3 for transmission.

At a remote ISU, CXSU performs the same steps 4420 for upstream



messages. Step 4422 is required only in an MISU, because all channels in an HISU reside in the same subband. Method 4420 may be performed at both ends of the system, as described herein; it may also be performed only at one end, either HDT or ISU. Staggering from the ISU to the HDT is preferable if only one end is staggered, because the most critical processing load is the error-correction of all 240 channels in the upstream receiving modem, Figure 26.

The use of error-correcting codes along with unencoded data raises problems in a real-time transport system. Data arrives from the trunk line 20, Figure 1, at a constant rate. This data must be transmitted downstream in the same time duration, whether it is encoded along the way, or sent unencoded. Likewise, upstream data usually must be transmitted at the same rate whether or not it is encoded. That is, the use of error-correcting codes must be time-transparent at both ends of the system. But error-correcting codes require the transmission of check digits or symbols along with the data. The present system resolves this difficulty by packing the data words differently if they are encoded. As explained above, the basic unencoded word length for a DS0++ channel is ten bits: eight data bits, a signaling (NBS) bit, and a parity bit. When encoding is used, however, this format is changed to nine-bit words, with a single parity bit for the entire message. This is the reason for the choice of frame sizes for the encoded modes. A 21-frame message contains 19 data frames, which would ordinarily be transmitted as  $10 \times 19 = 190$  bits. Those same data frames, packaged as nine-bit words along with two nine-bit check words, require  $(19+2) \times 9 = 189$  bits; adding one more parity bit covering the entire message lock gives 190 bits, the same number as that required for the unencoded version of the same data. The 41-frame message has 37 frames of data, or 370 bits in unencoded 10-bit format. Encoded as 37 nine-bit words along with four check words, the same message requires  $(37+4) \times 9 = 369$  bits; again, a single additional parity bit yield the 370 bits of the same data in unencoded form. The 81-frame format has 73 data words, 8 check words, and a parity bit, yielding the same number of bits as 73 data words in 10-bit form.

There are many other combinations of numbers which yield similar results. These can be found heuristically without a great deal of experimentation. The first step is to estimate rough numbers of large (parity-bearing) words in one or more message sizes, and the number of errors desired to be correctable for each size. The next step is to determine a number of smaller (non-parity) words that carry the same amount of data, but which form a total message the same size or slightly smaller than the total number of bits in the large-word format. Any excess bits then are assigned to parity over the block -- or to any other function, for that matter. For example, if two bits are left over instead of one for each message, they could represent two parity bits over the message, two control or format-designation bits, etc. The use of check symbols, of course, greatly reduces the need for parity or other forms of error detection. In fact, while the present system uses the message-parity bit as parity in the downstream direction, the ISU deliberately sets the parity bit to an incorrect value in an upstream message if it was incorrect in the downstream message. This serves to signal the HDT that a bit error was encountered, when the HDT would not otherwise be aware of it; this in turn allows the HDT to keep more accurate statistics on channel quality for reallocating channels, or for other purposes.

Figure 57 shows a method 4430 for adding the "code packing" feature to the method 4420 of Figure 56. Step 4431 repeats the steps for all channels. Step 4432 determines whether the data for the channel is to be encoded or not. If not, step 4433 merely transmits it word-by-word to the modem. If it is to be encoded, step 4434 strips the parity (or other) bit(s) from each word. After step 4435 has formed the check words, step 4436 calculates the message-wide parity, or other desired function. Thereafter, step 4437 waits the proper number of frame times (as specified by method 4420, and step 4438 sends the message to the modem as before.

In the upstream direction, the reverse path through the HDT is substantially a mirror of the forward path through the HDT 12. For example, the tenth parity bit is processed at the CXMU 56 and the signal from the CXMU 56 to the CTSU 54 is in

the format of Figure 9.

The round trip delay of a DSO is the same for every data path. The time delay over the path from the downstream CTSU output, through CXMU 56, over the HFC distribution network to the ISU 100 and then from the ISU 100, back over the HFC distribution network 11, through CXMU 56 and to CTSU 54 is controlled by upstream synchronization, as described in detail below. Generally, path delay is measured for each ISU and if it is not the correct number of frames long, the delay length is adjusted by adding delay to the path at the ISU 100.

#### Coax Master Unit (CXMU)

The coax master unit 56 (CXMU), shown in Figure 3, includes the coax master card logic 80 (CXMC) and the master coax card (MCC) modem 82. As previously described, up to six CXMUs may be equipped in an HDT 12. The 6 CXMUs 56 include three pairs of CXMUs 56 with each pair providing for transmit in a 6 MHz bandwidth. Each pair of CXMUs 56 includes one active CXMU and a standby CXMU. Thus, one to one protection for each CXMU is provided. As shown in Figure 3, both CXMUs of the pair are provided with upstream telephony data from the upstream telephony receiver 16 and are capable of transmitting via the coaxial line 22 to the downstream telephony transmitter 14. As such, only a control signal is required to provide for the one-to-one protection indicating which CXMU 56 of the pair is to be used for transmission or reception.

#### Coax Master Card Logic (CXMC)

The coax master card logic 80 (CXMC) of the CXMU 56 (Figure 12), provides the interface between the data signals of the HDT 12, in particular of the CTSU 54, and the modem interface for transport of data over the HFC distribution network 11. The CXMC 80 interfaces directly to the MCC modem 82. The CXMC 80 also implements an ISU operations channel transceiver for multi-point to point operation between the HDT 12 and all ISUs 100 serviced in the 6 MHz bandwidth

in which the CXMU 56 controls transport of data within. Referring to Figure 12, the CXMC includes controller and logic 84, downstream data conversion 88, upstream data conversion 90, data integrity 92, IOC transceiver 96, and timing generator 94.

Downstream data conversion 88 performs the conversion from the nine-bit channel format from CTSU 54 (Figure 9) to the ten-bit channel format (Figure 10) and generates the data integrity bit in each downstream channel transported over the HFC distribution network 11. The data integrity bit represents odd parity. Downstream data conversion 88 is comprised of at least a FIFO buffer used to remove the 32 gap bits 72, 74 (Figure 9) present in the downstream CTSU outputs and insert the tenth, data integrity bit, on each channel under control of controller and logic 84.

The upstream data conversion 90 includes at least a FIFO buffer which evaluates the tenth bit (data integrity) appended to each of the upstream channels and passes this information to the data integrity circuitry 92. The upstream data conversion 90 converts the data stream of ten-bit channels (Figure 10) back to the nine-bit channel format (Figure 9) for application to CTSU 54. Such conversion is performed under control of controller and logic 84.

The controller and logic 84 also manages call processing and channel allocation for the telephony transport over the HFC distribution network 11 and maintains traffic statistics over the HFC distribution network 11 in modes where dynamic time-slot allocation is utilized, such as for providing TR-303 services, concentration services commonly known to those skilled in the art. In addition, the controller 84 maintains error statistics for the channels in the 6 MHz band in which the CXMU transports data, provides software protocol for all ISU operations channel communications, and provides control for the corresponding MCC modem 82.

The data integrity 92 circuitry processes the output of the tenth bit evaluation of each upstream channel by the upstream conversion circuit 90. In the present system, parity is only guaranteed to be valid on a provisioned channel which has a

call in progress. Because initialized and activated ISU transmitters may be powered down when the ISUs are idle, the parity evaluation performed by the CXMC is not always valid. A parity error detected indicates either a transmission error in an upstream channel or a transmission error in a downstream channel corresponding to the upstream channel.

The ISU operations channel (IOC) transceiver 96 of the CXMC 80 contains transmit buffers to hold messages or control data from the controller and logic 84 and loads these IOC control messages which are a fixed total of 8 bytes in length into a 64 kbps channel to be provided to the MCC modem 82 for transport on the HFC distribution network 11. In this implementation, all IOC channels carry the same information at all times. That is, the IOC messages are broadcast simultaneously over all the channels. This allows the use of inexpensive and rugged narrow-band modems in the ISUs, reserving the more expensive and critical wideband models for the HDT, which uses only one modem for an entire 6MHz band, and which can be located centrally in a controlled environment. In the upstream direction, the IOC transceiver receives the 64 kbps channel via the MCC modem 82 which provides the controller and logic 84 with such messages.

The timing generator circuit 94 receives redundant system clock inputs from both the active and protection CTSUs 54 of the HDT 12. Such clocks include a 2 kHz HFC multiframe signal, which is generated by the CTSU 54 to synchronize the round trip delay on all the coaxial legs of the HFC distribution network. This signal indicates multiframe alignment on the ISU operations channel and is used to synchronize symbol timing and data reconstruction for the transport system. A 8 kHz frame signal is provided for indicating the first "gap" bit of a 2.56 MHz, 32 channel signal from the CTSU 54 to the CXMU 56. A 2.048 MHz clock is generated by the CTSU 54 to the SCNU 58 and the CXMU 56. The CXMU 56 uses this clock for ISU operations channel and modem communication between the CXMC 80 and the MCC modem 82. A 2.56 MHz bit clock is used for transfer of data signals between the DS1Us 48 and CTSUs 54 and the CTSUs 54 and CXMCs

56. A 20.48 MHz bit clock is utilized for transfer of the 10-bit data channels between the CXMC and the MCC.

#### Master Coax Card (MCC) Modem

5       The master coax card (MCC) modem 82 of the CXMU 56 interfaces on one side to the CXMC 80 and on the other side to the telephony transmitter 14 and receiver 16 for transmission on and reception from the HFC distribution network 11. The MCC modem 82 implements the modem functionality for OFDM transport of telephony data and control data. The block diagram of Figure 3 identifies the associated interconnects of the MCC modem 82 for both upstream and downstream communication. The MCC modem 82 is not an independent module in the HDT 12, as it has no interface to the HDT 12 other than through the CXMC 80 of the CXMU 56. The MCC modem 82 represents the transport system logic of the HDT 12. As such, it is responsible for implementing all requirements associated with information transport over the HFC distribution network 11. Each MCC modem 82 of the CXMUs 56 of HDT 12 is allocated a maximum bandwidth of 6 MHz in the downstream spectrum for telephony data and control data transport. The exact location of the 6 MHz band is provisionable by the CXMC 80 over the communication interface via the IOC transceiver 96 between the CXMC 80 and MCC modem 82. The downstream transmission of telephony and control data is in the RF spectrum of about 725 to 800 MHz.

15       Each MCC modem 82 is allocated a maximum of 6 MHz in the upstream spectrum for receipt of control data and telephony data from the ISUs within the RF spectrum of about 5 to 40 MHz. Again, the exact location of the 6 MHz band is provisionable by the CXMC 80 over the communication interface between the CXMC 80 and the MCC modem 82.

20       The MCC modem 82 receives 256 DS0+ channels from the CXMC 80 in the form of a 20.48 MHz signal as described previously above. The MCC modem 82 transmits this information to all the ISUs 100 using the multicarrier modulation

technique based on OFDM as previously discussed herein. The MCC modem 82 also recovers 256 DS0+ multicarrier channels in the upstream transmission over the HFC distribution network and converts this information into a 20.48 Mbps stream which is passed to CXMC 80. As described previously, the multicarrier modulation technique involves encoding the telephony and control data, such as by quadrature amplitude modulation, into symbols, and then performing an inverse fast Fourier transform technique to modulate the telephony and control data on a set of orthogonal multicarriers.

Symbol alignment is a necessary requirement for the multicarrier modulation technique implemented by the MCC modem 82 and the ISU modems 101 in the ISUs 100. In the downstream direction of transmission, all information at an ISU 100 is generated by a single CXMU 56, so the symbols modulated on each multicarrier are automatically phase aligned. However, upstream symbol alignment at a receiver of the MCC modem 82 varies due to the multi-point to point nature of the HFC distribution network 11 and the unequal delay paths of the ISUs 100. In order to maximize receiver efficiency at the MCC modem 82, all upstream symbols must be aligned within a narrow phase margin. This is done by utilizing an adjustable delay parameter in each ISU 100 such that the symbol periods of all channels received upstream from the different ISUs 100 are aligned at the point they reach the HDT 12. This is part of the upstream synchronization process and shall be described further below. In addition, to maintain orthogonality of the multicarriers, the carrier frequencies used for the upstream transmission by the ISUs 100 must be frequency locked to the HDT 12.

Incoming downstream information from the CXMC 80 to the MCC modem 82 is frame aligned to the 2 kHz and 8 kHz clocks provided to the MCC modem 82. The 2 kHz multi-frame signal is used by the MCC modem 82 to convey downstream symbol timing to the ISUs as described in further detail below. This multiframe clock conveys the channel correspondence and indicates the multi-carrier frame structure so that the telephony data may be correctly reassembled at the ISU 100.

Two kHz represents the greatest common factor between 10 kHz (the modem symbol rate) and 8 kHz (the data frame rate).

All ISUs 100 will use the synchronization information inserted by the associated MCC modem 82 to recover all downstream timing required by the ISUs 100. This synchronization allows the ISUs 100 to demodulate the downstream information and modulate the upstream transmission in such a way that all ISU 100 transmissions received at the HDT 12 are synchronized to the same reference. Thus, the carrier frequencies used for all ISU 100 upstream transmission will be frequency locked to the HDT 12.

The symbol alignment is performed over synchronization channels in the downstream and upstream 6 MHz bandwidths under the responsibility of the MCC modem 82, in addition to providing path delay adjustment, initialization and activation, and provisioning over such synchronization channels until initialization and activation is complete as further described herein. These parameters are then tracked by use of the IOC channels. Because of their importance in the system, the IOC channel and synchronization channels may use a different modulation scheme for transport of control data between the MCC modem 82 and ISUs 100 which is more robust or of lesser order (less bits/sec/Hz or bits/symbol) than used for transport of telephony data. For example, the telephony data may be modulated using quadrature amplitude modulation, while the IOC channel and synchronization channel may be modulated utilizing BPSK modulation techniques.

The MCC modem 82 also demodulates telephony and control data modulated on multicarriers by the ISUs 100. Such demodulation is described further below with respect to the various embodiments of the telephony transport system.

Functions with respect to the OFDM transport system for which the MCC modem 82 is responsible, include at least the following, which are further described with respect to the various embodiments in further detail. The MCC modem 82 detects a received amplitude/level of a synchronization pulse/pattern from an ISU 100 within a synchronization channel and passes an indication of this level to the



CXMC 80 over the communication interface therebetween. The CXMC 80 then provides a command to the MCC modem 82 for transmission to the ISU 100 being leveled for adjustment of the amplitude level thereof. The MCC modem 82 also provides for symbol alignment of all the upstream multicarriers by correlating an upstream pattern modulated on a synchronization channel with respect to a known symbol boundary and passing a required symbol delay correction to the CXMC 80 over the communication interface therebetween. The CXMC 80 then transmits via the MCC modem 82 a message downstream to the ISU 100 to adjust the symbol delay of the ISU 100.

Likewise, with regard to synchronizing an ISU 100 for overall path delay adjustment, the MCC modem 82 correlates an upstream multiframe pattern modulated in the proper bandwidth by the ISU 100 on the IOC channel with respect to a known reference boundary, and passes a required path delay correction to the CXMC 80 over the modem interface therebetween. The CXMC 80 then transmits via the MCC modem 82 over the IOC channel a message downstream to adjust the overall path delay of an ISU 100.

#### Summary of Bidirectional Multi-Point to Point Telephony Transport

The following summarizes the transport of telephony and control information over the HFC distribution network 11. Each CXMU 56 of HDT 12 is provisioned with respect to its specific upstream and downstream operating frequencies. The bandwidth of both upstream and downstream transmission by the CXMU 56 are a maximum of 6 MHz, with the downstream transmission in a 6 MHz band of the RF spectrum of about 725-800 MHz.

In the downstream direction, each MCC modem 82 of the CXMU 56 provides electrical telephony and control data signals to the downstream telephony transmitter 14 via coaxial line 22 in its provisional 6 MHz bandwidth. The RF electrical telephony and control data signals from the MCC modems 82 of the HDT 12 are combined into a composite signal. The downstream telephony transmitter

then passes the combined electrical signal to redundant electrical-to-optical converters for modulation onto a pair of protected downstream optical feeder lines 24.

The downstream optical feeder lines 24 carry the telephony information and control data to an ODN 18. At the ODN 18, the optical signal is converted back to electrical and combined with the downstream video information (from the video head-end feeder line 42) into an electrical downstream RF output signal. The electrical RF output signal including the telephony information and control data is then fed to the four coaxial distribution legs 30 by ODN 18. All telephony information and control data downstream is broadcast on each coaxial leg 30 and carried over the coaxial portion of the HFC distribution network 11. The electrical downstream output RF signal is tapped from the coax and terminated on the receiver modem 101 of an ISU 100 through duplex filter 104, shown in Figure 8.

The RF electrical output signals include telephony information and control data modulated on orthogonal multicarriers by MCC modem 82 utilizing orthogonal frequency division multiplexing techniques; the telephony information and control data being mapped into symbol data and the symbols being modulated on a plurality of orthogonal carriers using fast Fourier transform techniques. As the symbols are all modulated on carriers at a single point to be transmitted to multiple points in the system 10, orthogonality of the multicarriers and symbol alignment of the symbols modulated on the orthogonal multicarriers are automatically aligned for transport over the HFC distribution network 11 and the telephony information and control data is demodulated at the ISUs 100 by the modem 101.

The ISU 100 receives the RF signal tapped from the coax of the coaxial portion of the HFC network 11. The RF modem 101 of the ISU 100 demodulates the signal and passes the telephony information and control data extracted to the CXSU controller 102 for provision to channel units 103 as appropriate. The ISU 100 represents the interface where the telephony information is converted for use by a subscriber or customer.

The CXMUs 56 of the HDT 12 and the ISUs 100 implement the bidirectional multi-point to point telephony transport system of the communication system 10. The CXMUs 56 and the ISUs, therefore, carry out the modem functionality. The transport system in accordance with the present invention may utilize three different modems to implement the modem functionality for the transport system. The first modem is the MCC modem 82 which is located in each CXMU 56 of the HDT 12. The HDT 12, for example, includes three active MCC modems 82 (Figure 3) and is capable of supporting many ISUs 100, representing a multi-point to point transport network. The MCC modem 82 coordinates telephony information transport as well as control data transport for controlling the ISUs 100 by the HDT 12. For example, the control data may include call processing messages, dynamic allocation and assignment messages, ISU synchronization control messages, ISU modem control messages, channel unit provisioning, and any other ISU operation, administration, maintenance and provisioning (OAM&P) information.

The second modem is a single family subscriber or HISU modem optimized to support a single dwelling residential unit. Therefore, it must be low in cost and low in power consumption. The third modem is the multiple subscriber or MISU modem, which is required to generally support both residential and business services.

The HISU modem and the MISU modem may take several forms. For example, the HISU modem and the MISU modem may, as described further in detail below with regard to the various embodiments of the present invention, extract only a small portion of the multicarriers transmitted from the HDT 12 or a larger portion of the multicarriers transmitted from the HDT 12. For example, the HISU may extract 20 multicarriers or 10 payload channels of telephony information transported from the HDT 12 and the MISU may extract information from 260 multicarriers or 130 payload channels transported from the HDT 12. Each of these modems may use a separate receiver portion for extracting the control data from the signal transported by the HDT 12 and an additional receiver portion of the HISU

modem to extract the telephony information modulated on the multicarriers transported from the HDT 12. This shall be referred to hereinafter as an out of band ISU modem. The MCC modem 82 for use with an out of band ISU modem may modulate control information within the orthogonal carrier waveform or on carriers somewhat offset from such orthogonal carriers. In contrast to the out of band ISU modem, the HISU and MISU modems may utilize a single receiver for the ISU modem and extract both the telephony information and control data utilizing the single receiver of the modem. This shall be referred to hereinafter as an in-band ISU modem. In such a case, the control data is modulated on carriers within the orthogonal carrier waveform but may utilize different carrier modulation techniques. For example, BPSK for modulation of control data on the carriers as opposed to modulation of telephony data on payload carriers by QAM techniques. In addition, different modulation techniques may be used for upstream or downstream transmission for both control data and telephony data. For example, downstream telephony data may be modulated on the carriers utilizing 256 QAM and upstream telephony data may be modulated on the carriers utilizing 32 QAM. Whatever modulation technique is utilized for transmission dictates what demodulation approach would be used at the receiving end of the transport system. Demodulation of the downstream telephony information and control data transported by the HDT 12 shall be explained in further detail below with reference to block diagrams of different modem embodiments.

In the upstream direction, each ISU modem 101 at an ISU 100 transmits upstream on at least one orthogonal multicarrier in a 6 MHz bandwidth in the RF spectrum of about 5 to 40 MHz; the upstream 6 MHz band corresponding to the downstream 6 MHz band in which transmissions are received. The upstream electrical telephony and control data signals are transported by the ISU modems 101 to the respectively connected optical distribution node 18 as shown in Figure 1 via the individual coaxial cable legs 30. At the ODN 18, the upstream signals from the various ISUs are combined and transmitted optically to the HDT 12 via optical

feeder lines 26. As previously discussed, the upstream electrical signals from the various ISUs may, in part, be frequency shifted prior to being combined into a composite upstream optical signal. In such a case, the telephony receiver 16 would include corresponding downshifting circuitry.

5           Due to the multi-point to point nature of transport over the HFC distribution network 11 from multiple ISUs 100 to a single HDT 12, in order to utilize orthogonal frequency division multiplexing techniques, symbols modulated on each carrier by the ISUs 100 must be aligned within a certain phase margin. In addition, as discussed in further detail below, the round trip path delay from the network  
10       interface 62 of the HDT 12 to all ISUs 100 and back from the ISUs 100 to the network interface 62 in the communication system 10 must be equal. This is required so that signaling multiframe integrity is preserved throughout the system. In addition, a signal of proper amplitude must be received at the HDT 12 to perform any control functions with respect to the ISU 100. Likewise, with regard to OFDM  
15       transport from the ISUs 100, the ISUs 100 must be frequency locked to the HDT 12 such that the multicarriers transported over the HFC distribution network 11 are orthogonally aligned. The transport system implements a distributed loop technique for implementing this multi-point to point transport utilizing orthogonal frequency division multiplexing as further described below. When the HDT 12 receives the  
20       plurality of multicarriers which are orthogonally aligned and which have telephony and control data modulated thereon with symbols aligned, the MCC modems 82 of the CXMUs 56 demodulate the telephony information and control data from the plurality of multicarriers in their corresponding 6 MHz bandwidth and provide such telephony data to the CTSU 54 for delivery to the network interface 62 and the  
25       control data to the CXMC 80 for control of the telephony transport.

As one skilled in the art will recognize, the spectrum allocations, frequency assignments, data rates, channel numbers, types of services provided and any other parameters or characteristics of the system which may be a choice of design are to be taken as examples only. The invention as described in the accompanying claims

contemplates such design choices and they therefore fall within the scope of such claims. In addition, many functions may be implemented by software or hardware and either implementation is contemplated in accordance with the scope of the claims even though reference may only be made to implementation by one or the other.

#### First Embodiment of Telephony Transport System

The first embodiment of the telephony transport system in accordance with the present invention shall be described with particular reference to Figures 13-35 which include block diagrams of MCC modems 82, and HISU modems and MISU modems shown generally as ISU modem 101 in Figure 8. Such modems implement the upstream and downstream modem transport functionality. Following this description is a discussion on the theory of operation utilizing such modems.

Referring to Figure 13, the spectrum allocation for one 6 MHz band for upstream and downstream transport of telephony information and control data utilizing OFDM techniques is shown. The waveform preferably has 240 payload channels or DS0+ channels which include 480 carriers or tones for accommodating a net data rate of 19.2 Mbps, 24 IOC channels including 46 carriers or tones, and 2 synchronization channels. Each synchronization channel includes two carriers or tones and is each offset from 24 IOC channels and 240 payload channels by 10 unused carriers or tones, utilized as guard tones. The total carriers or tones is 552. The synchronization tones utilized for synchronization functions as described further below are located at the ends of the 6 MHz spectrum and the plurality of orthogonal carriers in the 6 MHz band are separated from carriers of adjacent 6 MHz bands by guard bands (516.0 kHz) at each end of the 6 MHz spectrum. The guard bands are provided at each end of the 6 MHz band to allow for filter selectivity at the transmitter and receivers of the system. The synchronization carriers are offset from the telephony data or payload carriers such that if the synchronization carrier utilized for synchronization during initialization and activation is not orthogonal

with the other tones or carriers within the 6 MHz band, the synchronization signal is prevented from destroying the structure of the orthogonally aligned waveform. The synchronization tones are, therefore, outside of the main body of payload carriers of the band and interspersed IOC channels, although the synchronization channel could be considered a special IOC channel.

To minimize the power requirement of the ISUs, the amount of bandwidth that an ISU processes is minimized. As such, the telephony payload channels and IOC channels of the 6 MHz band are interspersed in the telephony payload channels with an IOC channel located every 10 payload channels. With such a distributed technique, wherein subbands of payload channels greater than 10 include an IOC channel, the amount of bandwidth an ISU "sees" can be limited such that an IOC channel is available for the HDT 12 to communicate with the ISU 100. Such subband distribution for the spectral allocation shown in Figure 13 is shown in Figure 16. There are 24 subbands in the 6 MHz bandwidth with each subband including 10 payload channels with an IOC channel between the 5th and 6th payload channels. A benefit of distributing the IOC channels throughout the 6 MHz band is protection from narrow band ingress. If ingress destroys an IOC channel, there are other IOC channels available and the HDT 12 can re-tune an ISU 100 to a different portion of the 6 MHz band, where an IOC channel that is not corrupted is located.

Preferably, the MISU 66 sees approximately 3 MHz of the 6 MHz bandwidth to receive up to 130 payload channels which bandwidth also includes numerous IOC channels for communication from the HDT 12 to the MISU 66. The HISU 68 sees about 100 kHz of the 6 MHz bandwidth to receive 11 channels including at least one IOC channel for communication with the HDT 12.

The primary difference between the downstream and upstream paths are the support of downstream synchronization and upstream synchronization. In the downstream direction, all ISUs lock to information from the HDT (point to multi-point). The initialization and activation of ISUs are based on signals supplied in the

upstream synchronization channel. During operation, ISUs track the synchronization via the IOC channels. In the upstream, the upstream synchronization process involves the distributed (multi-point to point) control of amplitude, frequency, and timing; although frequency control can also be provided utilizing only the downstream synchronization channel as described further below. The process of upstream synchronization occurs in one of the two upstream synchronization channels, the primary or the secondary synchronization channel.

Referring to Figure 21, the downstream transmission architecture of the MCC modem 82 is shown. Two serial data inputs, approximately 10 Mbps each, comprise the payload data from the CXMC 56 which is clocked by the 8 kHz frame clock input. The IOC control data input from the CXMC 56 is clocked by the IOC clock input, which is preferably a 2.0 kHz clock. The telephony payload data and the IOC control data enter through serial ports 132 and the data is scrambled as known to one skilled in the art by scrambler 134 to provide randomness in the waveform to be transmitted over the HFC distribution network 11. Without scrambling, very high peaks in the waveform may occur; however, if the waveform is scrambled the symbols generated by the MCC modem 82 become sufficiently random and such peaks are sufficiently limited.

Figure 58 details the operation of a typical scrambler, such as 134, Figure 21. Symbol clock 4501 clocks a seed pattern through a linear-feedback shift register 4510 having nine stages, 4510-0 through 4510-8. With XOR gate 4511 positioned as shown, the generator polynomial is binary "100 010 000". The seed initially loaded into register 4510 at input 4502 is "111 001 110". Two identical translation tables 4520 and 4521 receive two-bit inputs from register 4510 at every symbol time. The high- and low-order bits of table 4520 proceed from the outputs of stages 4510-7 and 4510-6, respectively. High-order bit 4523 of table 4521 also receives output 4510-6, but as its high-order bit; stage output 4510-5 provides its low-order bit. Logic gates 4530 perform an XOR between the five-bit output of table 4520 and the upper five bits of a 10-bit DS0 word, while gates 4531 do the same for the lower



five bits of the same DS0 word. Outputs 4505 and 4506 carry the two 5-bit scrambled symbols for the DS0 word. Each descrambler such as 176, Figure 22 or 23, is identical to its corresponding scrambler. It recovers the original bit pattern of each symbol by decoding it with the same polynomial and seed.

5           The polynomial and seed for register 4510 of the scramblers and descramblers selected by known techniques to yield a maximal-length pseudo-random sequence. Inversion of the order of the input bits as between table 4520 and table 4521 increases the scrambling of the two symbols of the DS0 word. To increase the randomness among different sequences even more, different scramblers in the system have different polynomials and seeds. Randomness could be further increased by using more than four different table entries; however, the added complexity overrode the gain, for this particular embodiment. Only the payload channels are scrambled; the IOC channels are not scrambled.

10           The scrambled signals are applied to a symbol mapping function 136. The symbol mapping function 136 takes the input bits and maps them into a complex constellation point. For example, if the input bits are mapped into a symbol for output of a BPSK signal, every bit would be mapped to a single symbol in the constellation as in the mapping diagram for BPSK of Figure 15. Such mapping results in in-phase and quadrature values (I/Q values) for the data. BPSK is the modulation technique preferably used for the upstream and downstream IOC channels and the synchronization channels. BPSK encoding is preferred for the IOC control data so as to provide robustness in the system as previously discussed. For QPSK modulation, every two bits would map into one of four complex values that represent a constellation point. In the preferred embodiment, 32 QAM is utilized for telephony payload data, wherein every five bits of payload data is mapped into one of 32 constellation points as shown in Figure 14. Such mapping also results in I/Q values. As such, one DS0+ signal (10 bits) is represented by two symbols and the two symbols are transmitted using two carriers. Thus, one DS0+ channel is transported over two carriers or tones of 6 MHz spectrum.

One skilled in the art will recognize that various mapping or encoding techniques may be utilized with different carriers. For example, telephony channels carrying ISDN may be encoded using QPSK as opposed to telephony channels carrying POTS data being encoded using 32 QAM. Therefore, different telephony channels carrying different services may be modulated differently to provide for more robust telephony channels for those services that require such quality. The architecture in accordance with the present invention provides the flexibility to encode and modulate any of the channels differently from the modulation technique used for a different channel.

Within the framework of QAM32 modulation, Figure 17 shows a constellation which has improved characteristics. Here, the in-phase and quadrature values are shown encoded by three bits each instead of the four shown in Figure 13; their analog values, however, they remain in the ranges -5 to +5. The constellation of Figure 17 approaches as closely as possible to an analogy to a Gray code scheme, in which a transition from one row to the next and from one column to the next result in only a single bit change in the 5-bit symbol code. (The exceptions are four transitions from the first column to the second, and from the fifth to the sixth, which have two transitions each. The corner cells have zero transitions between these columns, which do not detract from the advantages of the scheme.) If a symbol is received incorrectly after transmission, the most likely error is a slight change in either amplitude or phase. If the bit strings represented by the symbols have as few bit transitions as possible for single-value phase and amplitude changes, then a reception error will create fewer bit errors on the final digital output. That is, small (symbol) errors in produce small (bit) errors out.

The constellations shown in Figures 14 and 17 use all points of a six-cell square except the four corners. Hence, they have two axes of symmetry, and appear identical when rotated by 90°, 180°, and 270°. If a phase error ever exceeds 45°, an attempted correction may pull the phase to an incorrect orientation. This is called four-fold phase ambiguity. However, deliberately using one and only one of the

corner points as a valid symbol provides a key for identifying the correct phase for errors as great as a full  $180^\circ$ . For example, designating the symbol for "16" as  $I=010 (+5)$  and  $Q=010 (+5)$  instead of the  $I=001, Q=010 (+3, +5)$  in Figure 17 introduces a symbol at this corner point whenever a "16" is sent upstream or downstream. Because only one corner is used, any received value having both I and Q values  $\pm 5$  requires phase rotation until  $I=+5$  and  $Q=+5$ . This assignment also preserves the nearly Gray-code structure of the constellation.

Any other symbol assignment which breaks the symmetry of the constellation would produce the same effect. Even a constellation retaining only one axis of symmetry would allow twice the phase-correction range of the constellation of Figure 17. For example, using both the upper left and lower right corners as valid symbols allows correction of phase errors up to  $90^\circ$ .

Each symbol that gets represented by the I/Q values is mapped into a fast Fourier transform (FFT) bin of symbol buffer 138. For example, for a DS0+, running at 8 kHz frame rate, five bits are mapped into one FFT bin and five bits into another bin. Each bin or memory location of the symbol buffer 138 represents the payload data and control data in the frequency domain as I/Q values. One set of FFT bins gets mapped into the time domain through the inverse FFT 140, as is known to one skilled in the art. The inverse FFT 140 maps the complex I/Q values into time domain samples corresponding to the number of points in the FFT. Both the payload data and IOC data are mapped into the buffer 138 and transformed into time domain samples by the inverse FFT 140. The number of points in the inverse FFT 140 may vary, but in the preferred embodiment the number of points is 256. The output of the inverse FFT 140, for a 256 point FFT, is 256 time domain samples of the waveform.

In conventional practice, buffer 138 clocks symbols into inverse FFT 140 at exactly the same rate that inverse FFT 140 clocks out the in-phase and quadrature values FFT I and FFT Q in Figure 21. To put the matter another way, the 256 digital waveform samples from buffer 138 represent  $360^\circ$ , or  $2\pi$  radians, of a QAM

32 waveform having the amplitude and phase of the 5 bits of its symbol, as determined by mapping unit 136. The FFT I and Q outputs represent 256 samples of a frequency spectrum corresponding to the same time period. At the receiving end, however, any misalignment at all in the phase synchronization causes FFT 170, Figure 22, or 180, Figure 23, to decode a portion of a previous or subsequent symbol's waveform along with somewhat less than the full cycle of the desired symbol; this intersymbol interference can cause misreading the symbol as a different valid symbol, resulting in as many as five bit errors.

In a presently preferred embodiment, the 256 samples clocked into inverse FFT 140 represent an extra  $45^\circ$  ( $\pi/4$  radians) above a complete cycle. Another way to think of this is that the symbols are clocked into the FFT at an effective 9 kHz rate, and clocked out at the nominal 8 kHz symbol rate. Figure 52 shows an unmodulated sine wave, i.e., one having  $I=0$ ,  $Q=0$  in the units used herein. The upper portion shows one cycle,  $0-360^\circ$ , at the nominal 8 kHz frame rate. The lower portion shows the same wave at a 9 kHz rate, so that the amount of time previously occupied by  $360^\circ$  now takes up  $405^\circ$  of phase — from  $-22.5^\circ$  to  $+382.5^\circ$ . Obviously, there are phase discontinuities between successive cycles of the wave. Figure 53 shows a typical QAM 32 wave modulated at a different amplitude and a slightly different phase from those of Figure 52. These might correspond to, say,  $I=-1$ ,  $Q=+1$  in the scheme used herein. The small portions at the ends of this wave represent unmodulated cycles, as in Figure 52. The phase of this wave is advanced from the corresponding wave of the lower portion of Figure 52; it does not cross the zero axis at  $0^\circ$  and  $180^\circ$  of its proper cycle. It does, however, include the extra  $22.5^\circ$  of excess phase at each end, for  $45^\circ$  extra over an 8 kHz cycle. Again, a phase discontinuities exist at the ends of the total  $405^\circ$  phase degrees of this wave.

In fact, this characteristic gives the excess-phase improvement an advantage over its primary function of providing a guard band for the symbol decoder, for reducing intersymbol interference. In Figure 52, successive cycles of a wave modulated with the same symbol (or with no symbol), produce a continuous

waveform with no breaks or other features to distinguish the beginnings and endings of individual cycles. The lower part of this diagram demonstrates that even an unmodulated excess-phase waveform contains discontinuity features serving as markers at the ends of each cycle. A repeating string of idle symbols, or any other symbols, likewise produces these markers. In the frequency- and phase-acquisition and tracking aspects discussed below, such markers therefore provide definite waveform features for synchronizing purposes, without having to guarantee the transmission of any special string of varying characters strictly for synchronization. This saves the overhead of interrupting the payload and/or IOC channels to provide such a string, and the complexity of storing or diverting payload information while the sync string is present. It also allows sync to take place at times when, because of the above factors, it would not be feasible otherwise.

At the receiving end, FFT 170 (in an MISU) or 180 (HISU) decodes the 256 time slots for one frame time as  $405^\circ$  of a cycle to symbol decoder 174, which matches the cycle to the nearest 5-bit string of bits. Because any phase difference up to  $\pm 22.5^\circ$  will never conflate the proper wave with that for another symbol, no intersymbol interference at all occurs within this margin of error in phase tracking. This provides a form of guard band for each symbol. In the upstream direction, units 186, 188, and 190 or 191 provide excess phase in the transmitting MISU and HISU modems of Figures 24 and 25; and the head-end receiving modem of Figure 26 decodes and tracks this phase as described above.

The inverse FFT 140 has separate serial outputs for in-phase and quadrature (I/Q) components, FFT1 and FFT0. Digital to analog converters 142 take the in-phase and quadrature components, which is a numeric representation of baseband modulated signal and convert it to a discrete waveform. The signal then passes through reconstruction filters 144 to remove harmonic content. This reconstruction is needed to avoid problems arising from multiple mixing schemes and other filtering problems. The signal is summed in a signal conversion transmitter 146 for up-converting the I/Q components utilizing a synthesized waveform that is digitally

tunable with the in-phase and quadrature components for mixing to the applicable transmit frequency. For example, if the synthesizer is at 600 MHz, the output frequency will be at 600 MHz. The components are summed by the signal conversion transmitter 146 and the waveform including a plurality of orthogonal carriers is then amplified by transmitter amplifier 148 and filtered by transmitter filter 150 before being coupled onto the optical fiber by way of telephony transmitter 14. Such functions are performed under control of general purpose processor 149 and other processing circuitry of block 147 necessary to perform such modulation. The general purpose processor also receives ISU adjustment parameters from carrier, amplitude, timing recovery block 222 (Figure 26) for carrying out distributed loop symbol alignment, frequency locking, amplitude adjustment, and path delay functions as described further below.

In conventional practice, the relationship between the frequency of a carrier and the frequency and timing of data symbols modulated onto that carrier is arbitrary and unimportant. In the present system, however, it has been found that even very small frequency drifts between the 8 kHz symbol or frame clock and the frequencies of the tones upon which they ride can produce significant intersymbol interference and distortion at the receiving end. Such drifts tend to destroy the orthogonality of the channel signals produced by inverse FFT 140 in Figure 21. The present system also, however, provides a simple, inexpensive way to overcome this problem. Figure 51 shows a portion 4200 of the HDT clock/sync logic in CTSU 54, Figure 3. Timing recovery loop 4210 produces a single master reference clock output at 10.24 MHz. Although loop 4210 could be a free-running oscillator, it is in fact slaved to the network 10, Figure 1. With which the entire system communicates. This connection is convenient in eliminating gross or unpredictable differences between the data speeds of the network and the system.

Smoothing loop 4220 evens out short-term variations in the signal from loop 4210. Phase comparator 4221 controls a voltage-controlled crystal oscillator at

40.96 MHz; divider 4223 provides feedback at the proper frequency. Comparator 4221 includes a low-pass integrator which gives phase-lock loop 4220 a bandwidth of about 130Hz. Divider 4230 reduces the frequency of VCXO 4222 to 2.56Mhz. A second phase-lock loop 4240 has a phase comparator 4241, again with low-pass characteristics, feeding a voltage-controlled oscillator running at 1267.2MHz; divider provides feedback at the proper frequency. Divider 4250 produces the final RF clock frequency, 9.9MHz, at output 4251. The network clock is sufficiently accurate over long periods of time, but it is subject to significant amounts of short-period jitter. The large amount of smoothing provided by loops 4220 and 4240 overcome the intolerance of analog RF components for short-term variations.

Meanwhile, digital divider 4260 divides the master 10.24MHz clock by a factor of 80 to produce an 8kHz symbol or frame clock output 4261. Output 4261 does not require the smoothing, because it clocks only digital circuits, which are relatively insensitive to short-term frequency changes.

RF master clock 4261 proceeds to RF synthesizer 143 in HDT transmitting modem 82, as shown in Figure 21, where it directly controls the frequency of the tunable 500-850MHz RF carrier for the entire band carrying all of the channels shown in Figures 13 and 16. Symbol clock 4261 proceeds to the frame-clock inputs in Figure 21, where it controls the symbol timing, and, because it also controls the FFT speed, the frequencies of the channels in the entire band. Clock lock 4200 thus provides a solid link which inherently preserves the orthogonality of the band signals in a multicarrier system, by deriving the RF carrier clock and the symbol or frame clock from the same source. At the same time, it provides a small amount of gradual variation for satisfying the demands of the analog RF components.

The overall purpose of locking the two clocks together at the HDT is to lock the carrier clocks and the symbol (frame) clocks throughout the system; and the purpose of this in turn is to preserve the orthogonality of the signals in a multicarrier system which is capable of bidirectional operation: that is, as a multipoint-to-point-configuration as well as in the usual point-to-multipoint "broadcast" direction.

Clock generator 166, Figures 22 and 23, of timing generator 107, Figure 6 locks to the frequencies of the incoming signals to provide the clocks used in the remote ISU modules. Therefore, the carrier and frame clocks in each upstream transmitter portion, Figure 24, of remote modem 108, Figure 8, are also locked to each other, by virtue of being locked to the incoming signal from the HDT.

At the downstream receiving end, either an MISU or an HISU provides for extracting telephony information and control data from the downstream transmission in one of the 6 MHz bandwidths. With respect to the MISU 66, the MISU downstream receiver architecture is shown in Figure 22. It includes a 100 MHz bandpass filter 152 to reduce the frequency band of the received 600 to 850 MHz total band broadcast downstream. The filtered signal then passes through voltage tuned filters 154 to remove out of band interference and further reduce the bandwidth. The signal is down converted to baseband frequency via quadrature and in-phase down converter 158 where the signal is mixed at complex mixers 156 utilizing synthesizer 157 which is controlled from an output of serial ports 178. The down converted I/Q components are passed through filters 159 and converted to digital format at analog to digital converters 160. The time domain samples of the I/Q components are placed in a sample buffer 162 and a set of samples are input to down converter compensation unit 164. The compensation unit 164 attempts to mitigate errors such as DC offsets from the mixers and differential phase delays that occur in the down conversion.

Carrier, amplitude and timing signaling are extracted from the compensated signal, by the carrier, amplitude, and timing recovery block 166 by extracting control data from the synchronization channels during initialization and activation of the ISU and the IOC channels during tracking as further described below with reference to Figure 33. The compensated signal in parallel form is provided to fast Fourier transform (FFT) 170 to be converted into a vector of frequency domain elements which are essentially the complex constellation points with I/Q components originally created upstream at the MCC modem 82 for the DS0+



channels which the MISU sees. Due to inaccuracies in channel filtering, an equalizer 172 removes dynamic errors that occur during transmission and reception. Equalization in the upstream receiver and the downstream receiver architectures shall be explained in further detail below with reference to Figure 35. From the equalizer 172, the complex constellation points are converted to bits by symbol to bit converter 174, descrambled at descrambler 176 which is a mirror element of scrambler 134, and the payload telephony information and IOC control data are output by the serial ports 178 to the CXSU 102 as shown in Figure 8. Block 153 includes the processing capabilities for carrying out the various functions as shown therein.

Referring to Figure 23, the HISU 68 downstream receiver architecture is shown. The primary difference between the HISU downstream receiver architecture (Figure 23) and the MISU downstream receiver architecture (Figure 22) is the amount of bandwidth being processed. The front ends of the receivers up to the FFT processing are substantially the same, except during the down conversion, the analog to digital converters 160 can be operated at a much slower rate. For instance, if the bandwidth of the signal being processed is 100 kHz, the sample rate can be approximately 200 kHz. In an MISU processing a 3 MHz signal, the sample rate is about 6 MHz. Since the HISU is limited to receiving a maximum of 10 DS0+s, the FFT 180 can be of a smaller size. A 32 point FFT 180 is preferably used in the HISU and can be implemented more efficiently, compared to a 128 or 256 point FFT utilized in the MISU. Therefore, the major difference between these architectures is that the HISU receiver architecture requires substantially less signal processing capability than the MISU receiver and as such has less power consumption. Thus, to provide a system wherein power consumption at the remote units is minimized, the smaller band of frequencies seen by the HISU allows for such low consumption. One reason the HISU is allowed to see such a small band of carriers is that the IOC channels are interspersed throughout the 6 MHz spectrum.

Referring to Figure 24, the upstream transmission architecture for the HISU

68 is shown. The IOC control data and the telephony payload data from the CXSU 102 (Figure 8) is provided to serial ports 182 at a much slower rate in the HISU than in the MISU or HDT transmission architectures, because the HISU supports only 10 DS0+ channels. The HISU upstream transmission architecture implements three important operations. It adjusts the amplitude of the signal transmitted, the timing delay (both symbol and path delay) of the signal transmitted, and the carrier frequency of the signal transmitted. The telephony data and IOC control data enters through the serial ports 182 under control of clocking signals generated by the clock generator 173 of the HISU downstream receiver architecture, and is scrambled by scrambler 184 for the reasons stated above with regard to the MCC downstream transmission architecture. The incoming bits are mapped into symbols, or complex constellation points, including I/Q components in the frequency domain, by bits to symbol converter 186. The constellation points are then placed in symbol buffer 188. Following the buffer 188, an inverse FFT 190 is applied to the symbols to create time domain samples; 32 samples corresponding to the 32 point FFT. A delay buffer 192 is placed on the output of the inverse FFT 190 to provide multi-frame alignment at MCC modem upstream receiver architecture as a function of the upstream synchronization process controlled by the HDT 12. The delay buffer 192, therefore, provides a path delay adjustment prior to digital to analog conversion by the digital to analog converters 194 of the in-phase and quadrature components of the output of the inverse FFT 190. Clock delay 196 provides a fine tune adjustment for the symbol alignment at the request of IOC control data output obtained by extracting control data from the serial stream of data prior to being scrambled. After conversion to analog components by digital to analog converters 194, the analog components therefrom are reconstructed into a smooth analog waveform by the reconstruction filters 198. The upstream signal is then directly up converted by direct converter 197 to the appropriate transmit frequency under control of synthesizer block 195. Synthesizer block 195 is operated under control of commands from an IOC control channel which provides carrier frequency

00001344-00001

adjustment commands thereto as extracted in the HISU downstream receiver architecture. The up converted signal is then amplified by transmitter amplifier 200, filtered by transmitter filter 202 and transmitted upstream to be combined with other signals transmitted by other ISUs 100. The block 181 includes processing circuitry for carrying out the functions thereof.

Referring to Figure 27, the upstream transmitter architecture for the MISU 66 is shown and is substantially the same as the upstream transmitter architecture of HISU 68. However, the MISU 66 handles more channels and cannot perform the operation on a single processor as can the HISU 68. Therefore, both a processor of block 181 providing the functions of block 181 including the inverse FFT 190 and a general purpose processor 206 to support the architecture are needed to handle the increased channel capacity.

Referring to Figure 26, the MCC upstream receiver architecture of each CXMU 56 at the HDT 12 is shown. A 5 to 40 MHz band pass filter 208 filters the upstream signal which is then subjected to a direct down conversion to baseband by mixer and synthesizer circuitry 211. The outputs of the down conversion is applied to anti-alias filters 210 for conditioning thereof and the output signal is converted to digital format by analog to digital converters 212 to provide a time domain sampling of the in-phase and quadrature components of the signal to narrow band ingress filter and FFT 112. The narrow band ingress filter and FFT 112, as described below, provides protection against narrow band interference that may affect the upstream transmission.

The ingress filter and FFT 112 protects ten channels at a time, therefore, if ingress affects one of the available 240 DS0+s in the 6 MHz spectrum received by MCC modem 82, a maximum of ten channels will be destroyed from the ingress. The ingress filter and FFT 112 includes a polyphase structure, as will be recognized by one skilled in the art as a common filter technique. It will be further recognized by one skilled in the art that the number of channels protected by the polyphase filter can be varied. The output of the ingress filter and FFT 112 is coupled to an

equalizer 214 which provides correction for inaccuracies that occur in the channel, such as those due to noise from reference oscillators or synthesizers. The output symbols of the equalizer 214, are applied to a symbols to bits converter 216 where the symbols are mapped into bits. The bits are provided to descramblers 218, which are a mirror of the scramblers of the ISUs 100 and the output of the descramblers are provided to serial ports 220. The output of the serial ports is broken into two payload streams and one IOC control data stream just as is provided to the MCC downstream transmitter architecture in the downstream direction. Block 217 includes the necessary processing circuitry for carrying out the functions therein.

In order to detect the downstream information, the amplitude, frequency, and timing of the arriving signal must be acquired using the downstream synchronization process. Since the downstream signal constitutes a point to multi-point node topology, the OFDM waveform arrives via a single path in an inherently synchronous manner, in contrast to the upstream signal. Acquisition of the waveform parameters is initially performed on the downstream synchronization channels in the downstream synchronization bands located at the ends of the 6 MHz spectrum. These synchronization bands include a single synchronization carrier or tone which is BPSK modulated by a 2 kHz framing clock. This tone is used to derive initial amplitude, frequency, and timing at the ISU. The synchronization carrier may be located in the center of the receive band and could be considered a special case of an IOC. After the signal is received and the receiver architecture is tuned to a typical IOC channel, the same circuitry is used to track the synchronization parameters using the IOC channel. The process used to acquire the necessary signal parameters utilizes carrier, amplitude and timing recovery block 166 of the ISU receiver architecture, which is shown in more detail in block diagram form in Figure 33. The carrier, amplitude and timing recovery block 166 includes a Costas loop 330 which is used to acquire the frequency lock for the received waveform. After the signal is received from the compensation unit 164, a sample and hold 334 and analog to digital conversion 332 is applied to the signal with the

resulting samples from the converters 332 applied to the Costas loop 330. The sampling is performed under control of voltage controlled oscillator 340 as divided by divider 333 which divides by the number of points of the FFT utilized in the receiver architecture, M. The mixers 331 of the Costas loop 330 are fed by the arriving signal and the feedback path, and serve as the loop phase detectors. The output of the mixers 331 are filtered and decimated to reduce the processing requirements of subsequent hardware. Given that the received signal is band-limited, less samples are required to represent the synchronization signal. If orthogonality is not preserved in the receiver, the filter will eliminate undesired signal components from the recovery process. Under conditions of orthogonality, the LPF 337 will completely remove effects from adjacent OFDM carriers. When carrier frequency lock is achieved, the process will reveal the desired BPSK waveform in the in-phase arm of the loop. The output of the decimators are fed through another mixer, then processed through the loop filter with filter function  $H(s)$  and numerically controlled oscillator (NCO), completing the feedback path to correct for frequency error. When the error is at a "small" level, the loop is locked. In order to achieve fast acquisition and minimal jitter during tracking, it will be necessary to employ dual loop bandwidths. System operation will require that frequency lock is achieved and maintained within about  $\pm 4\%$  of the OFDM channel spacing (360 Hz).

The amplitude of the signal is measured at the output of the frequency recovery loop at BPSK power detector 336. The total signal power will be measured, and can be used to adjust a numerically controllable analog gain circuit (not shown). The gain circuit is intended to normalize the signal so that the analog to digital converters are used in an optimal operating region.

Timing recovery is performed using an early-late gate type algorithm of early-late gate phase detector 338 to derive timing error, and by adjusting the sample clock or oscillator 340 in response to the error signal. The early-late gate detector results in an advance/retard command during an update interval. This command

will be applied to the sample clock or oscillator 340 through filter 341. This loop is held off until frequency lock and amplitude lock have been achieved. When the timing loop is locked, it generates a lock indicator signal. The same clocks are also used for the upstream transmission. The carrier, timing and amplitude recovery block 166 provides a reference for the clock generator 168. The clock generator 168 provides all of the clocks needed by the MISU, for example, the 8 kHz frame clock and the sample clock.

Carrier, amplitude, and timing recovery block 222 of the MCC modem upstream receiver architecture (Figure 26), is shown by the synchronization loop diagram of Figure 34. It performs detection for upstream synchronization on signals on the upstream synchronization channel. For initialization and activation of an ISU, upstream synchronization is performed by the HDT commanding one of the ISUs via the downstream IOC control channels to send a reference signal upstream on a synchronization channel. The carrier, amplitude, and timing recovery block 222 measures the parameters of data from the ISU 100 that responds on the synchronization channel and estimates the frequency error, the amplitude error, and the timing error compared to references at the HDT 12. The output of the carrier, amplitude, and timing recovery block 222 is turned into adjustment commands by the HDT 12 and sent to the ISU being initialized and activated in the downstream direction on an IOC control channel by the MCC downstream transmitter architecture.

The purpose of the upstream synchronization process is to initialize and activate ISUs such that the waveform from distinct ISUs combine to a unified waveform at the HDT 12. The parameters that are estimated at the HDT 12 by carrier, amplitude, and timing recovery block 222 and adjusted by the ISUs are amplitude, timing, and frequency. The amplitude of an ISUs signal is normalized so that DS0+s are apportioned an equal amount of power, and achieves a desired signal to noise ratio at the HDT 12. In addition, adjacent ISUs must be received at the correct relative level or else weaker DS0+ channels will be adversely impacted by

5

10

15

25

difference with respect to the time will be used to estimate the frequency error. The frequency error is generated by filter function  $H(s)$  444 and provided to the CXMC 80 for processing to send a frequency adjustment command to the ISU via an IOC control channel. The frequency error is also applied to the numerically controlled oscillator (NCO) to complete the frequency loop to correct for frequency error.

The amplitude error is computed based on the magnitude of the carrier during the upstream synchronization by detecting the carrier amplitude of the in-phase arm of the Costas loop 430 by power detector 436. The amplitude is compared with a desired reference value at reference comparator 443 and the error will be sent to the CXMC 80 for processing to send an amplitude adjustment command to the ISU via an IOC control channel.

When the local reference in the HDT has achieved phase lock, the BPSK signal on the synchronization channel arriving from the ISU is available for processing. The square wave is obtained on the in-phase arm of the Costas loop 430 and applied to early-late gate phase detector 438 for comparison to the locally generated 8 kHz signal from divider 435. The phase detector 438 generates a phase or symbol timing error applied to loop filter 441 and output via line 439. The phase or symbol timing error is then provided to the CXMC 80 for processing to send a symbol timing adjustment command to the ISU via an IOC control channel.

The mechanisms in the ISU which adjust the parameters for upstream synchronization include implementing an amplitude change with a scalar multiplication of the time domain waveform as it is being collected from the digital processing algorithm, such as inverse FFT 190, by the digital to analog converters 194 (Figure 24). Similarly, a complex mixing signal could be created and implemented as a complex multiply applied to the input to the digital to analog converters 194.

Frequency accuracy of both the downstream sample clock and upstream sample clock, in the ISU, is established by phase locking an oscillator to the downstream synchronization and IOC information. Upstream transmission



frequency is adjusted, for example, at synthesizer block 195 as commanded by the HDT 12.

Symbol timing corrections are implemented as a delay function. Symbol timing alignment in the ISU upstream direction is therefore established as a delay in the sample timing accomplished by either blanking a sample interval (two of the same samples to go out simultaneously) or by putting in an extra clock edge (one sample is clocked out and lost) via clock delay 196 (Figure 24). In this manner, a delay function can be controlled without data storage overhead beyond that already required.

After the ISU is initialized and activated into the system, ready for transmission, the ISU will maintain required upstream synchronization system parameters using the carrier, amplitude, frequency recovery block 222. An unused but initialized and activated ISU will be commanded to transmit on an IOC and the block 222 will estimate the parameters therefrom as explained above.

In both the upstream transmitter architectures for the MISU 66 (Figure 24) and the HISU 68 (Figure 25), frequency offset or correction to achieve orthogonality of the carriers at HDT 12 can be determined on the ISU as opposed to the frequency offset being determined at the HDT during synchronization by carrier, amplitude and timing recovery block 222 (Figure 26) and then frequency offset adjustment commands being transmitted to the ISU for adjustment of carrier frequency via the synthesizer blocks 195 and 199 of the HISU 68 and MISU 66, respectively. Thus, frequency error would no longer be detected by carrier, amplitude and timing recovery block 222 as described above. Rather, in such a direct ISU implementation, the ISU, whether an HISU 68 or MISU 66, estimates a frequency error digitally from the downstream signal and a correction is applied to the upstream data being transmitted.

The HDT 12 derives all transmit and receive frequencies from the same fundamental oscillator. Therefore, all mixing signals are frequency locked in the HDT. Similarly, the ISU, whether an HISU 68 or MISU 66, derives all transmit and

receive frequencies from the same fundamental oscillator; therefore, all the mixing signals on the ISU are also frequency locked. There is, however, a frequency offset present in the ISU oscillators relative to the HDT oscillators. The amount of frequency error (viewed from the ISU) will be a fixed percentage of the mixing frequency. For example, if the ISU oscillator is 10PPM off in frequency from the HDT oscillators, and the downstream ISU receiver mix frequency was 100 MHz and the ISU upstream transmit mixing frequency were 10 MHz, the ISU would have to correct for 1 kHz on the downstream receiver and create a signal with a 100 Hz offset on the upstream transmitter. As such, with the ISU direct implementation, the frequency offset is estimated from the downstream signal.

The estimation is performed with digital circuitry performing numeric calculations, i.e. a processor. Samples of the synchronization channel or IOC channel are collected in hardware during operation of the system. A tracking loop drives a digital numeric oscillator which is digitally mixed against the received signal. This process derives a signal internally that is essentially locked to the HDT. The internal numerical mix accounts for the frequency offset. During the process of locking to the downstream signal in the ISU, the estimate of frequency error is derived and with the downstream frequency being known, a fractional frequency error can be computed. Based on the knowledge of the mixing frequency at the HDT that will be used to down convert the upstream receive signal, an offset to the ISU transmit frequency is computed. This frequency offset is digitally applied to the ISU transmitted signal prior to converting the signal to the analog domain, such as by converters 194 of Figure 24. Therefore, the frequency correction can be performed directly on the ISU.

Referring to Figures 31 and 32, the narrow band ingress filter and FFT 112 of the MCC upstream receiver architecture, including a polyphase filter structure, will be described in further detail. Generally, the polyphase filter structure includes polyphase filters 122 and 124 and provides protection against ingress. The 6 MHz band of upstream OFDM carriers from the ISUs 100 is broken into subbands

through the polyphase filters which provide filtering for small groups of carriers or tones, and if an ingress affects carriers within a group of carriers, only that group of carriers is affected and the other groups of carriers are protected by such filtering characteristics.

5           The ingress filter structure has two parallel banks 122, 124 of polyphase filters. One bank has approximately 17 different non-overlapping bands with channel spaces between the bands. A magnitude response of a single polyphase filter bank is shown in Figure 29. The second bank is offset from the first bank by an amount so that the channels that are not filtered by the first bank are filtered by  
10           the second bank. Therefore, as shown in the closeup magnitude response of a single polyphase filter bank in Figure 30, one band of channels filtered may include those in frequency bins 38-68 with the center carriers corresponding to bins 45-61 being passed by the filter. The overlapping filter provides for filtering carriers in the spaces between the bands and the carriers not passed by the other filter bank. For  
15           example, the overlapping filter may pass 28-44. The two channel banks are offset by 16 frequency bins so that the combination of the two filter banks receives every one of the 544 channels.

Referring to Figure 31, the ingress filter structure receives the sampled  
20           waveform  $x(k)$  from the analog to digital converters 212 and then complex mixers 118 and 120 provide the stagger for application to the polyphase filters 122, 124. The mixer 118 uses a constant value and the mixer 120 uses a value to achieve such offset. The outputs of each mixer enters one of the polyphase filters 122, 124. The output of each polyphase filter bank comprises 18 bands, each of which contain 16 usable FFT bins or each band supports sixteen carriers at the 8 kHz rate, or 8 DS0+s.  
25           One band is not utilized.

Each band output of the polyphase filters 122, 124 has 36 samples per 8 kHz frame including 4 guard samples and enters a Fast Fourier Transform (FFT) block 126, 128. The first operation performed by the FFT blocks 126, 128 is to remove the four guard samples, thereby leaving 32 time domain points. The output of the

each FFT in the blocks is 32 frequency bins, 16 of which are used with the other bins providing filtering. The output of the FFTs are staggered to provide overlap. As seen in Figure 31, carriers 0 - 15 are output by FFT #1 of the top bank, carriers 16 - 31 are output by FFT #1 of the bottom bank, carriers 32 - 48 are output by FFT #2 of the top bank and so on.

The polyphase filters 122, 124 are each standard polyphase filter construction as is known to one skilled in the art and each is shown by the structure of Figure 32. The input signal is sampled at a 5.184 mega- sample per second rate, or 648 samples per frame. The input is then decimated by a factor of 18 (1 of 18 samples are kept) to give an effective sample rate of 288 kHz. This signal is subjected to the finite impulse response (FIR) filters, labeled  $H_{0,0}(Z)$  through  $H_{0,16}(Z)$ , which include a number of taps, preferably 5 taps per filter. As one skilled in the art will recognize the number of taps can vary and is not intended to limit the scope of the invention. The outputs from the filters enter an 18 point inverse FFT 130. The output of the transform is 36 samples for an 8 kHz frame including 4 guard samples and is provided to FFT blocks 126 and 128 for processing as described above. The FFT tones are preferably spaced at 9 kHz, and the information rate is 8 kilosymbols per second with four guard samples per symbol allotted. The 17 bands from each polyphase filter are applied to the FFT blocks 126, 128 for processing and output of the 544 carriers as indicated above. One band, the 18th band, as indicated above, is not used.

The equalizer 214 (Figure 26) and 172 (Figure 22), in both upstream and downstream receiver architectures, is supplied to account for changes in group delay across the cable plant. The equalizer tracks out phase and gain or amplitude variations due to environmental changes and can therefore adapt slowly while maintaining sufficiently accurate tracking. The coefficients 360 of the equalizer 172, 214, for which the internal equalizer operation is generally shown in Figure 35, represent the inverse of the channel frequency response to the resolution of the FFT 112, 170. The downstream coefficients will be highly correlated since every

channel will progress through the same signal path as opposed to the upstream coefficients which may be uncorrelated due to the variant channels that individual DS0+s may encounter in the multi-point to point topology. While the channel characteristics are diverse, the equalizer will operate the same for either upstream or downstream receiver.

The downstream equalizer will track on only the IOC channels, thus reducing the computational requirements at the ISUs and removing the requirement for a preamble in the payload channels, as described further below, since the IOC channels are always transmitted. The upstream, however, will require equalization on a per DS0+ and IOC channel basis.

The algorithm used to update the equalizer coefficients contains several local minima when operating on a 32 QAM constellation and suffers from a four-fold phase ambiguity. Furthermore, each DS0+ in the upstream can emanate from a separate ISU, and can therefore have an independent phase shift. To mitigate this problem, each communication onset will be required to post a fixed symbol preamble prior to data transmission. Note that the IOC channels are excluded from this requirement since they are not equalized and that the preamble cannot be scrambled. It is known that at the time of transmission, the HDT 12 will still have accurate frequency lock and symbol timing as established during initialization and activation of the ISU and will maintain synchronization on the continuously available downstream IOC channel.

The introduction of the preamble requires that the equalizer have knowledge of its process state. Three states are introduced which include: search, acquisition, and tracking mode. Search mode is based on the amount of power present on a channel. Transmitter algorithms will place a zero value in unused FFT bins, resulting in no power being transmitted on that particular frequency. At the receiver, the equalizer will determine that it is in search mode based on the absence of power in the FFT bin.

When transmission begins for an initialized and activated ISU, the equalizer detects the presence of signal and enters the acquisition mode. The length of the preamble may be about 15 symbols. The equalizer will vary the equalization process based on the preamble. The initial phase and amplitude correction will be large but subsequent updates of the coefficients will be less significant. In order to differentiate the training pattern from any other data sequence, when the HDT informs an ISU to connect a new payload channel, the ISU transmits 16 consecutive symbols having  $I=0$  and  $Q=0$ , which is not a valid data symbol in the constellations of Figures 14 or 17. The ISU then transmits 8 valid data symbols, allowing the equalizer for that channel to set its coefficient properly to adjust for amplitude and phase of the incoming signal.

After acquisition, the equalizer will enter a tracking mode with the update rate being reduced to a minimal level. The tracking mode will continue until a loss of power is detected on the channel for a period of time. The channel is then in the unused but initialized and activated state. The equalizer will not train or track when the receiver is being tuned and the coefficients will not be updated. The coefficients may be accessed and used such as by signal to noise detector 305 (Figure 26) for channel monitoring as discussed further below.

For the equalization process, the I/Q components are loaded into a buffer at the output of the FFT, such as FFT 112, 180. As will be apparent to one skilled in the art, the following description of the equalizer structure is with regard to the upstream receiver equalizer 214 but is equally applicable to the downstream receiver equalizer 172. The equalizer 214 extracts time domain samples from the buffer and processes one complex sample at a time. The processed information is then output therefrom. Figure 35 shows the basic structure of the equalizer algorithm less the state control algorithm which should be apparent to one skilled in the art. The primary equalization path performs a complex multiply at multiplier 370 with the value from the selected FFT bin. The output is then quantized at symbol quantize block 366 to the nearest symbol value from a storage table. The quantized value

(hard decision) is passed out to be decoded into bits by symbols to bits converter 216. The remainder of the circuitry is used to update the equalizer coefficients. An error is calculated between the quantized symbol value and the equalized sample at summer 364. This complex error is multiplied with the received sample at multiplier 363 and the result is scaled by the adaptation coefficient by multiplier 362 to form an update value. The update value is summed at summer 368 with the original coefficient to result in a new coefficient value.

#### Operation of First Embodiment

In the preferred embodiment, the 6 MHz frequency band for each MCC modem 82 of HDT 12 is allocated as shown in Figure 13. Although the MCC modem 82 transmits and receives the entire 6 MHz band, the ISU modems 100 (Figure 8) are optimized for the specific application for which they are designed and may terminate/generate fewer than the total number of carriers or tones allocated in the 6 MHz band. The upstream and downstream band allocations are preferably symmetric. The upstream 6 MHz bands from the MCC modems 82 lie in the 5-40 MHz spectrum and the downstream 6 MHz bands lie in the 725-760 MHz spectrum. One skilled in the art will recognize that if different transmission media are utilized for upstream and downstream transmission, the frequencies for transmission may be the same or overlap but still be non-interfering.

There are three regions in each 6 MHz frequency band to support specific operations, such as transport of telephony payload data, transport of ISU system operations and control data (IOC control data), and upstream and downstream synchronization. Each carrier or tone in the OFDM frequency band consists of a sinusoid which is modulated in amplitude and phase to form a complex constellation point as previously described. The fundamental symbol rate of the OFDM waveform is 8 kHz, and there are a total of 552 tones in the 6 MHz band. The following Table 2 summarizes the preferable modulation type and bandwidth allocation for the various tone classifications.

Table 2

Band Allocation	Number of Tones or Carriers	Modulation	Capacity	Bandwidth
Synch Band	24 tones(2 synch tones at each end and 10 guard tones at each end)	BPSK	n/a	216 kHz
Payload Data	480 (240 DS0+ channels)	32 QAM	19.2 MBPS	4.32 MHz
IOC	48 (2 every 20 data channels or 24 IOC channels)	BPSK	384 kbps	432 kHz
Intra-band guard	Remainder on each end	n/a	n/a	1.032 MHz (516 kHz at each end)
Composite Signal	552	n/a	n/a	6.0 MHz

Guard bands are provided at each end of the spectrum to allow for selectivity filtering after transmission and prior to reception. A total of 240 telephony data channels are included throughout the band, which accommodates a net data rate of 19.2 Mbps. This capacity was designed to account for additive ingress, thereby retaining enough support to achieve concentration of users to the central office. The IOC channels are interspersed throughout the band to provide redundancy and communication support to narrowband receivers located in the HISUs. The IOC data rate is 16 kbps (two BPSK tones at the symbol rate of 8 kHz frames per second). Effectively, an IOC is provided for every 10 payload data channels. An ISU, such as an HISU, that can only see a single IOC channel would be forced to retune if the IOC channel is corrupted. However, an ISU which can see multiple



IOC channels can select an alternate IOC channel in the event that the primary choice is corrupt, such as for an MISU.

The synchronization channels are duplicated at the ends of the band for redundancy, and are offset from the main body of usable carriers to guarantee that the synchronization channels do not interfere with the other used channels. The synchronization channels were previously described and will be further described below. The synchronization channels are operated at a lower power level than the telephony payload channels to also reduce the effect of any interference to such channels. This power reduction also allows for a smaller guard band to be used between the synchronization channels and the payload telephony channels.

One synchronization or redundant synchronization channels may also be implemented within the telephony channels as opposed to being offset therefrom. In order to keep them from interfering with the telephony channels, the synchronization channels may be implemented using a lower symbol rate. For example, if the telephony channels are implemented at an 8 kHz symbol rate, the synchronization channels could be implemented at a 2 kHz symbol rate and also may be at a lower power level.

The ISUs 100 are designed to receive a subband, as shown in Figure 16, of the total aggregate 6 MHz spectrum. As an example, the HISU 68 will preferably detect only 22 of the available 552 channels. This implementation is primarily a cost/power savings technique. By reducing the number of channels being received, the sample rate and associated processing requirements are dramatically reduced and can be achievable with common conversion parts on the market today.

A given HISU 68 is limited to receiving a maximum of 10 DS0s out of the payload data channels in the HISU receiver's frequency view. The remaining channels will be used as a guard interval. Furthermore, in order to reduce the power/cost requirements, synthesizing frequency steps will be limited to 198 kHz. An IOC channel is provided for as shown in Figure 16 so that every HISU 68 will always see an IOC channel for control of the HISU 68 via HDT 12.

The MISU 66 is designed to receive 13 subbands, as shown in Figure 16, or 130 of the 240 available DS0s. Again, the tuning steps will be limited to 198 kHz to realize an efficient synthesizer implementation. These are preferred values for the HISU 68 and MISU 66, and it will be noted by one skilled in the art that many of the values specified herein can be varied without changing the scope or spirit of the invention as defined by the accompanying claims.

As known to one skilled in the art, there may be need to support operation over channels in a bandwidth of less than 6 MHz. With appropriate software and hardware modifications of the system, such reconfiguration is possible as would be apparent to one skilled in the art. For example, for a 2 MHz system, in the downstream, the HDT 12 would generate the channels over a subset of the total band. The HISUs are inherently narrowband and would be able to tune into the 2 MHz band. The MISUs supporting 130 channels, would receive signals beyond the 2 MHz band. They would require reduction in filter selectivity by way of a hardware modification. An eighty (80) channel MISU would be able to operate with the constraints of the 2 MHz system. In the upstream, the HISUs would generate signals within the 2 MHz band and the MISUs transmit section would restrict the information generated to the narrower band. At the HDT, the ingress filtering would provide sufficient selectivity against out of band signal energy. The narrowband system would require synchronization bands at the edges of the 2 MHz band.

As previously described, acquisition of signal parameters for initializing the system for detection of the downstream information is performed using the downstream synchronization channels. The ISUs use the carrier, amplitude, timing recovery block 166 to establish the downstream synchronization of frequency, amplitude and timing for such detection of downstream information. The downstream signal constitutes a point to multi-point topology and the OFDM waveform arrives at the ISUs via a single path in an inherently synchronous manner. In the upstream direction, each ISU 100 must be initialized and activated through a

process of upstream synchronization before an HDT 12 can enable the ISU 100 for transmission. The process of upstream synchronization for the ISUs is utilized so that the waveform from distinct ISUs combine to a unified waveform at the HDT. The upstream synchronization process, portions of which were previously described, involves various steps. They include: ISU transmission level adjustment, upstream multicarrier symbol alignment, carrier frequency adjustment, and round trip path delay adjustment. Such synchronization is performed after acquisition of a 6 MHz band of operation.

Generally, with respect to level adjustment, the HDT 12 calibrates the measured signal strength of the upstream transmission received from an ISU 100 and adjusts the ISU 100 transmit level so that all ISUs are within acceptable threshold. Level adjustment is performed prior to symbol alignment and path delay adjustment to maximize the accuracy of these measurements.

Generally, symbol alignment is a necessary requirement for the multicarrier modulation approach implemented by the MCC modems 82 and the ISU modems 101. In the downstream direction of transmission, all information received at the ISU 100 is generated by a single CXMU 56, so the symbols modulated on each multicarrier are automatically phase aligned. However, upstream symbol alignment at the MCC modem 82 receiver architecture varies due to the multi-point to point nature of the HFC distribution network 11 and the unequal delay paths of the ISUs 100. In order to have maximum receiver efficiency, all upstream symbols must be aligned within a narrow phase margin. This is done by providing an adjustable delay path parameter in each ISU 100 such that the symbol periods of all channels received upstream from the different ISUs are aligned at the point they reach the HDT 12.

Generally, round trip path delay adjustment is performed such that the round trip delay from the HDT network interface 62 to all ISUs 100 and back to the network interface 62 from all the ISUs 100 in a system must be equal. This is required so that signaling multiframe integrity is preserved throughout the system.

All round trip processing for the telephony transport section has a predictable delay with the exception of the physical delay associated with signal propagation across the HFC distribution network 11 itself. ISUs 100 located at close physical distance from the HDT 12 will have less round trip delay than ISUs located at the maximum distance from the HDT 12. Path delay adjustment is implemented to force the transport system of all ISUs to have equal round trip propagation delay. This also maintains DS1 multiframe alignment for DS1 channels transported through the system, maintaining in-band channel signaling or robbed-bit signaling with the same alignment for voice services associated with the same DS1.

Generally, carrier frequency adjustment must be performed such that the spacing between carrier frequencies is such as to maintain orthogonality of the carriers. If the multicarriers are not received at the MCC modem 82 in orthogonal alignment, interference between the multicarriers may occur. Such carrier frequency adjustment can be performed in a manner like that of symbol timing or amplitude adjustment or may be implemented on the ISU as described previously above.

In the initialization process, when the ISU has just been powered up, the ISU 100 has no knowledge of which downstream 6 MHz frequency band it should be receiving in. This results in the need for the acquisition of 6 MHz band in the initialization process. Until an ISU 100 has successfully acquired a 6 MHz band for operation, it implements a "scanning" approach to locate its downstream frequency band. Generally, a local processor of the CXSU controller 102 of ISU 100 starts with a default 6 MHz receive frequency band somewhere in the range from 625 to 850 MHz. The ISU 100 waits for a period of time, for example 100 milliseconds, in each 6 MHz band to look for a valid 6 MHz acquisition command which matches a unique identification number for the ISU 100 after obtaining a valid synchronization signal; which unique identifier may take the form of or be based on a serial number of the ISU equipment. If a valid 6 MHz acquisition command or valid synchronization command is not found in that 6 MHz band, the CXSU controller 102 looks at the next 6 MHz band and the process is repeated. In this manner, as

explained further below, the HDT 12 can tell the ISU 100 which 6 MHz band it should use for frequency reception and later which band for frequency transmission upstream.

5 The process of initialization and activation of ISUs, as generally described above, and tracking or follow-up synchronization is further described below. This description is written using an MISU 66 in conjunction with a CXSU controller 102 but is equally applicable to any ISU 100 implemented with an equivalent controller logic. The coax master card logic (CXMC) 80 is instructed by the shelf controller unit (SCNU) 58 to initialize and activate a particular ISU 100. The SCNU uses an  
10 ISU designation number to address the ISU 100. The CXMC 80 correlates the ISU designation number with an equipment serial number, or unique identifier, for the equipment. No two ISU equipments shipped from the factory possess the same unique identifier. If the ISU 100 has never before been initialized and activated in the current system database, the CXMC 80 chooses a personal identification number (PIN) code for the ISU 100 being initialized and activated. This PIN code is then  
15 stored in the CXMC 80 and effectively represents the "address" for all communications with that ISU 100 which will follow. The CXMC 80 maintains a lookup table between each ISU designation number, the unique identifier for the ISU equipment, and the PIN code. Each ISU 100 associated with the CXMU 56 has  
20 a unique PIN address code assignment. One PIN address code will be reserved for a broadcast feature to all ISUs, which allows for the HDT to send messages to all initialized and activated ISUs 100.

The CXMC 80 sends an initialization and activation enabling message to the MCC modem 82 which notifies the MCC modem 82 that the process is beginning  
25 and the associated detection functionality in the MCC modem 82 should be enabled. Such functionality is performed at least in part by carrier, amplitude, timing recovery block 222 as shown in the MCC upstream receiver architecture of Figure 26 and as previously discussed.

The CXMC 80 transmits an identification message by the MCC modem 82

over all IOC channels of the 6 MHz band in which it transmits. The message includes a PIN address code to be assigned to the ISU being initialized and activated, a command indicating that ISU initialization and activation should be enabled at the ISU 100, the unique identifier for the ISU equipment, such as the equipment serial number, and cyclical redundancy checksum (CRC). The messages are sent periodically for a certain period of time,  $T_{SCAN}$ , which is shown as 6.16 seconds in Figure 20 and which is also shown in Figure 19. This period of time is the maximum time which an ISU can scan all downstream 6 MHz bands, synchronize, and listen for a valid identification message. The periodic rate, for example 50 msec, affects how quickly the ISU learns its identity. The CXMC 80 will never attempt to synchronize more than one ISU at a time, but will attempt to identify several ISUs during burst identification as described further below. A software timeout is implemented if an ISU does not respond after some maximum time limit is exceeded. This timeout must be beyond the maximum time limit required for an ISU to obtain synchronization functions.

During periodic transmission by CXMC 80, the ISU implements the scanning approach to locate its downstream frequency band. The local processor of the CXSU starts with a default 6 MHz receive frequency band somewhere in the range from 625 to 850 MHz. The ISU 100 selects the primary synchronization channel of the 6 MHz band and then tests for loss of synchronization after a period of time. If loss of synchronization is still present, the secondary synchronization channel is selected and tested for loss of synchronization after a period of time. If loss of synchronization is still present, then the ISU restarts selection of the synchronization channels on the next 6 MHz band which may be 1 MHz away but still 6 MHz in width. When loss of synchronization is not present on a synchronization channel then the ISU selects the first subband including the IOC and listens for a correct identification message. If a correct identification message is found which matches its unique identifier then the PIN address code is latched into an appropriate register. If a correct identification message is not found in the first subband on that IOC then

5

10

15

20

transmission in the upstream and the HDT 12 locks to the upstream signal. The HDT 12 derives an error indicator with regard to the parameter being adjusted and commands the subscriber ISU to adjust such parameter. The adjustment of error is repeated in the process until the parameter for ISU transmission is locked to the HDT 12.

More specifically, after the ISU 100 has acquired the 6 MHz band for operation, the CXSU 102 sends a message command to the ISU modem 101 and the ISU modem 101 transmits a synchronization pattern on a synchronization channel in the primary synchronization band of the spectral allocation as shown in Figures 13-18. The upstream synchronization channels which are offset from the payload data channels as allocated in Figures 13-18 include both a primary and a redundant synchronization channel such that upstream synchronization can still be accomplished if one of the synchronization channels is corrupted. The HDT monitors one channel for every ISU.

The MCC modem 82 detects a valid signal and performs an amplitude level measurement on a received signal from the ISU. The synchronization pattern indicates to the CXMC 80 that the ISU 100 has received the activation and initialization and frequency band commands and is ready to proceed with upstream synchronization. The amplitude level is compared to a desired reference level. The CXMC 80 determines whether or not the transmit level of the ISU 100 should be adjusted and the amount of such adjustment. If level adjustment is required, the CXMC 80 transmits a message on the downstream IOC channel instructing the CXSU 102 of the ISU 100 to adjust the power level of the transmitter of the ISU modem 101. The CXMC 80 continues to check the receive power level from the ISU 100 and provides adjustment commands to the ISU 100 until the level transmitted by the ISU 100 is acceptable. The amplitude is adjusted at the ISU as previously discussed. If amplitude equilibrium is not reached within a certain number of iterations of amplitude adjustment or if a signal presence is never detected utilizing the primary synchronization channel then the same process is used



on the redundant synchronization channel. If amplitude equilibrium is not reached within a certain number of iterations of amplitude adjustment or if a signal presence is never detected utilizing the primary or redundant synchronization channels then the ISU is reset.

5           Once transmission level adjustment of the ISU 100 is completed and has been stabilized, the CXMC 80 and MCC modem 82 perform carrier frequency locking. The MCC modem 82 detects the carrier frequency as transmitted by the ISU 100 and performs a correlation on the received transmission from the ISU to calculate a carrier frequency error correction necessary to orthogonally align the multicarriers of all the upstream transmissions from the ISUs. The MCC modem 82 returns a  
10           message to the CXMC 80 indicating the amount of carrier frequency error adjustment required to perform frequency alignment for the ISU. The CXMC 80 transmits a message on a downstream IOC channel via the MCC modem 82 instructing the CXSU 102 to adjust the transmit frequency of the ISU modem 101  
15           and the process is repeated until the frequency has been established to within a certain tolerance for the OFDM channel spacing. Such adjustment would be made via at least the synthesizer block 195 (Figures 24 and 25). If frequency locking and adjustment is accomplished on the ISU as previously described, then this frequency adjustment method is not utilized.

20           To establish orthogonality, the CXMC 80 and MCC modem 82 perform symbol alignment. The MCC modem 82 detects the synchronization channel modulated at a 8 kHz frame rate transmitted by the ISU modem 101 and performs a hardware correlation on the receive signal to calculate the delay correction necessary to symbol align the upstream ISU transmission from all the distinct ISUs 100. The  
25           MCC modem 82 returns a message to the CXMC 80 indicating the amount of delay adjustment required to symbol align the ISU 100 such that all the symbols are received at the HDT 12 simultaneously. The CXMC 80 transmits a message in a downstream IOC channel by the MCC modem 82 instructing the CXSU controller 102 to adjust the delay of the ISU modem 101 transmission and the process repeats

until ISU symbol alignment is achieved. Such symbol alignment would be adjusted via at least the clock delay 196 (Figures 24 and 25). Numerous iterations may be necessary to reach symbol alignment equilibrium and if it is not reached within a predetermined number of iterations, then the ISU may again be reset.

5 Simultaneously with symbol alignment, the CXMC 80 transmits a message to the MCC modem 82 to perform path delay adjustment. The CXMC 80 sends a message on a downstream IOC channel via the MCC modem 82 instructing the CXSU controller 102 to enable the ISU modem 101 to transmit a another signal on a synchronization channel which indicates the multiframe (2 kHz) alignment of the  
10 ISU 100. The MCC modem 82 detects this multiframe alignment pattern and performs a hardware correlation on the pattern. From this correlation, the modem 82 calculates the additional symbol periods required to meet the round trip path delay of the communication system. The MCC modem 82 then returns a message to the CXMC 80 indicating the additional amount of delay which must be added to  
15 meet the overall path delay requirements and the CXMC then transmits a message on a downstream IOC channel via the MCC modem 82 instructing the CXSU controller 102 to relay a message to the ISU modem 101 containing the path delay adjustment value. Numerous iterations may be necessary to reach path delay equilibrium and if it is not reached within a predetermined number of iterations, then the ISU may again be reset. Such adjustment is made in the ISU transmitter as  
20 can be seen in the display delay buffer "n" samples 192 of the upstream transmitter architectures of Figures 24 and 25. Path delay and symbol alignment may be performed at the same time, separately or together using the same or different signals sent on the synchronization channel.

25 Until the ISU is initialized and activated, the ISU 100 has no capability of transmitting telephony data information on any of the 480 tones or carriers. After such initialization and activation has been completed, the ISUs are within tolerance required for transmission within the OFDM waveform and the ISU is informed that transmission is possible and upstream synchronization is complete.

After an ISU 100 is initialized and activated for the system, follow-up synchronization or tracking may be performed periodically to keep the ISUs calibrated within the required tolerance of the OFDM transport requirements. The follow-up process is implemented to account for drift of component values with temperature. If an ISU 100 is inactive for extreme periods of time, the ISU can be tuned to the synchronization channels and requested to update upstream synchronization parameters in accordance with the upstream synchronization process described above. Alternatively, if an ISU has been used recently, the follow-up synchronization or tracking can proceed over an IOC channel. Under this scenario, as generally shown in Figure 28, the ISU 100 is requested to provide a signal over an IOC channel by the HDT 12. The HDT 12 then acquires and verifies that the signal is within the tolerance required for a channel within the OFDM waveform. If not, then the ISU is requested to adjust such errored parameters. In addition, during long periods of use, ISUs can also be requested by the HDT 12 to send a signal on an IOC channel or a synchronization channel for the purpose of updating the upstream synchronization parameters.

In the downstream direction, the IOC channels transport control information to the ISUs 100. The modulation format is preferably differentially encoded BPSK, although the differential aspect of the downstream modulation is not required. In the upstream direction, the IOC channels transport control information to the HDT 12. The IOC channels are differentially BPSK modulated to mitigate the transient time associated with the equalizer when sending data in the upstream direction. Control data is slotted on a byte boundary (500  $\mu$ s frame). Data from any ISU can be transmitted on an IOC channel asynchronously; therefore, there is the potential for collisions to occur.

As there is potential for collisions, detection of collisions on the upstream IOC channels is accomplished at a data protocol level. The protocol for handling such collisions may, for example, include exponential backoff by the ISUs. As such, when the HDT 12 detects an error in transmission, a retransmission command

is broadcast to all the ISUs such that the ISUs retransmit the upstream signal on the IOC channel after waiting a particular time; the wait time period being based on an exponential function.

One skilled in the art will recognize that upstream synchronization can be implemented allowing for multi-point to point transmission using only the symbol timing loop for adjustment of symbol timing by the ISUs as commanded by the HDT. The frequency loop for upstream synchronization can be eliminated with use of high quality local free running oscillators in the ISUs that are not locked to the HDT. In addition, the local oscillators at the ISUs could be locked to an outside reference. The amplitude loop is not essential to achieve symbol alignment at the HDT.

In the process described above with respect to initialization and activation, including upstream synchronization, if for some reason communication is lost between a large number of ISUs 100 and the HDT 12, after a period of time these ISUs 100 will require initialization and activation once again. Such a case may arise when a fiber is cut and users of multiple ISUs 100 are left without service. As initialization and activation is described above, only one ISU 100 would be initialized and activated at one time. The time frame for initialization and activation of multiple ISUs 100 in this manner is shown in Figure 19.

In Figure 19, each ISU 100 is initialized, as previously described, by identification of the ISU and acquisition by the ISU of the 6 MHz band for downstream transmission during a scanning period  $T_{SCAN}$ , which is the time period needed for the ISU to scan all of the downstream bands listening for a matching identification message. In one embodiment,  $T_{SCAN}$  is equal to 6.16 seconds. Of course this time period is going to be dependant on the number of bands scanned, the time period necessary for synchronizing on the downstream synchronization channels, and the time required to acquire an IOC channel within the band.

Further as shown in Figure 19, after each ISU has acquired a downstream and upstream 6 MHz band, upstream synchronization is then performed during a time

period  $T_{\text{EQUAL}}$ .  $T_{\text{EQUAL}}$  may be defined as the period in which an ISU should have received all messages from the CXMC 80 finishing the upstream synchronization process as described above, with a reasonable amount of iterations to accomplish such synchronization. At the very least, this time period is the time period necessary to accomplish symbol timing such that symbols received from various ISUs 100 at the HDT 12 are orthogonal. This time period would be increased if amplitude, frequency and path delay synchronization is also performed as described above. Therefore, the time period necessary to serially initialize and activate twelve ISUs,  $T_{\text{SERIAL}}$ , as shown in Figure 19 would be equal to  $12T_{\text{SCAN}} + 12T_{\text{EQUAL}}$ .

With a burst identification process as shall be described with respect to Figure 20, the time period for initializing and activating twelve ISUs 100 can be substantially reduced. This results in more ISUs 100 being activated more quickly and more users once again serviced in a shorter period of time. In the process of burst identification as shown by the timing diagram of Figure 20, the identification and acquisition of multiple ISUs 100 is performed in parallel instead of being performed serially as described above.

The periodicity of the identification messages sent by the CXMC 80 during initialization and acquisition, when performed during normal operating conditions when ISUs 100 are serially initialized, is designed to present a light load of traffic on the IOC channel but yet to allow a reasonable identification duration. This periodicity duration is, for example, 50 milliseconds. For the system to be able to handle both situations, serial identification and burst identification, this periodicity is kept the same. However, in burst identification, the IOC channel traffic load is not important because the service of all the ISUs 100 receiving communication via one of the CXMCs 80 utilizing the IOC channels has been terminated such as by the cut fiber. Therefore, during burst identification the IOC channels can be loaded more heavily and identification messages for multiple ISUs 100 utilizing such IOC channels are transmitted on the IOC channels at the same periodicity as during serial identification, but the phase for the identification messages is different for each ISU.

Due to the period, and utilization of the IOC channels for the identification messages during burst identification, there is a limit on the number of identification messages which may be transmitted during one  $T_{SCAN}$  period. If the periodicity is 50 milliseconds and the use of the IOC channel for a single identification message is 4 milliseconds, only twelve ISUs 100 may be identified during one  $T_{SCAN}$  period during burst identification. As described below with further reference to Figure 20, if the number of ISUs 100 desired to be burst identified is greater than twelve, then multiple groups of burst identifications are serially performed.

One skilled in the art will recognize that the specified numbers for time periods are for illustration only and that the present invention is not limited to such specified time periods. For example, the periodicity may be 100 milliseconds and the number of ISUs identified during burst identification may be 24. Many different time periods may be specified with corresponding changes made to the other times periods specified. Further, burst identification may be accomplished having a periodicity different than that for serial identification.

As shown by the timing diagram of Figure 19, a single burst initialization and activation of twelve inactive ISUs 100 can be accomplished in the time period  $T_{BURST}$  which is equal to  $T_{SCAN} + 12T_{EQUAL}$ . This is an  $11T_{SCAN}$  difference from the process carried out serially. During the  $T_{SCAN}$  period, identification messages for all twelve ISUs 100 being initialized are transmitted on the IOC channels for a CXMC 80. The twelve identification messages are each sent once during each 50 millisecond period. The phase of each message is however different. For example, identification message for ISU0 may be sent at time 0 and then again at 50 milliseconds, whereas the identification message for ISU1 may be sent at time 4 milliseconds and then again at 54 milliseconds and so forth.

After the ISUs 100 being initialized have been identified and acquisition of the downstream 6 MHz band has occurred during  $T_{SCAN}$ , then upstream synchronization is performed in a serial manner with respect to each ISU identified during  $T_{SCAN}$ . The upstream synchronization for the ISUs is performed during the

time period equal to  $12T_{\text{EQUAL}}$ . The CXMC 80 would start the upstream synchronization process in the same manner as described above for each ISU identified in a serial manner. The CXMC 80 sends to the ISU the upstream transmission band in which the ISU being synchronized is to transmit within and enables the upstream synchronization process to begin. The upstream synchronization process for an ISU has been described in detail above. If an upstream transmission band is not received and upstream synchronization is not enabled for an ISU during the  $12T_{\text{EQUAL}}$  time period, then the ISU is reset at the end of the  $12T_{\text{EQUAL}}$  period by a period of time equal to  $T_{\text{SCAN}} + 12T_{\text{EQUAL}}$  to possibly perform upstream synchronization in the next  $12T_{\text{EQUAL}}$  period. Once a burst identification period,  $T_{\text{BURST}}$ , is completed, the process may be started over again in a second  $T_{\text{BURST}}$  period as shown in Figure 20 if additional ISUs 100 are to be initialized and activated.

Figures 47 and 48 describe a control loop distributed through the system for acquiring and tracking an ISU, according to another embodiment of the invention. Loop 3900 shows steps executed by the ISU 66 or 68 at the left, and those executed by HDT 12 at the right. Messages between these two units are shown as horizontal dashed lines; the IOC channels carry these messages.

Steps 3910 prepare the ISU to communicate with the HDT. Step 3911 reads a prestored internal table 3912 indicating frequencies of the valid RF downstream bands, along with other information. Next, step 3913 tunes the ISU's narrow-band receive modem to the center of the first 6MHz band in table 3912. Step 3913a then fine-tunes it to one of the two sync channels in that band; assume that this "primary" channel is the one at the lower end of the downstream part of the band shown in Figure 13. Step 3914 acquires the amplitude and frequency of this sync tone. Briefly, equalizer 172, Figure 22 or 23, is adjusted to bring the out put of the FFT to about 12dB below its upper limit. Recovery block 166 measures the time for ten frames of the sync tone, and compares it to the 1.25msec correct time interval; the frequency of synthesizer 157 is then adjusted accordingly. Rough timing is

5

10

15-

20

25



messages 3933: The PIN is only one byte long, since addresses need be unique only within each 6MHz band, and are practical to use for identifying the ISU whenever the HDT needs to communicate with it. Message 3927 informs the HDT that the ISU will respond to its PIN.

5 Steps 3940 set up the upstream communications from the ISU to the HDT. After 3941 receives PIN confirmation 3927, step 3942 sends a designation of the upstream frequency band to the ISU as an IOC message 3943. This frequency may have been specified by the operator in step 3931, or may have been generated by the HDT itself. Step 3944 tunes the ISU modem to this 6Mhz band, and returns a  
10 confirmation 3945. Step 3946 then fine-tunes to the primary upstream sync channel of that band, such as the lower one in Figure 13. Step 3947 enables an HDT receiver on the designated band.

Steps 3950 adjust the transmitted power of the ISU in the upstream direction; in a multipoint-to-point-system, the power levels of all ISUs must track each other  
15 in order to ensure orthogonality of the entire signal received by the modem of Figure 26. Step 3951 transmits a ranging tone at an initial power level from the ISU on this sync channel, which is sometimes called a ranging channel. At the HDT, step 3952 measures the received power level, and block 3953 sends an IOC message 3954, causing step 3955 at the ISU to adjust the power of its transmitter 200,  
20 Figure 24, if necessary. If this cannot be done, message 3956 causes step 3946 to retune to a secondary ranging channel, such as the higher upstream sync tone of Figure 13, and causes the step 3947 to enable the secondary channel at the HDT. If this loop also fails, exit 3957 reports a hard failure to the system logic.

Blocks 3960 align the symbol or frame timing between the ISU and the HDT.  
25 Step 3961 measures the phase of the ISU's ranging tone with respect to the sync tone that the HDT is sending to all ISUs at all times. Figure 11 shows this signal, labeled "8kHz frame clk" in Figure 24. Step 3962 sends messages 3963 as necessary to cause the ISU modem to adjust the timing of its ranging tone, in step 3964. When this has completed, block 3965 detects whether or not groups of four

5

10

25

transmitting modem of Figure 21. The practicalities of a multicarrier (MC) system impose requirements which are not obvious from experience with TDMA (time division multiple-access) and other conventional forms of bidirectional multipoint networks and systems, nor in point-to-multipoint ("broadcast") multicarrier networks. In TDMA and similar systems, slight errors in frequency and timing, and larger errors in amplitude, can be compensated by tracking the received signal. In a broadcast MC system, all carriers are synchronized at the head end, and can again be tracked at the receiver. In a bidirectional multipoint-to-point multicarrier system, however, the HDT receiver must see all channels as though they had been generated by a single source, because the HDT decodes all channels in an entire 6MHz band as a single entity, with a single FFT. Even slight errors among the individual ISUs in their 10-channel and 130-channel subbands causes severe distortion and intersymbol interference when the HDT FFT inverts the channels back into symbol strings for multiple DS0 channels. The errors to be controlled are frequency, symbol timing, and signal amplitude. Amplitude (power level) in particular has been found to be much more stringent than in previous systems.

At all times after the ISU receiving modem Figure 22 or 23 has acquired the signal of HDT transmitting modem Figure 21, it must track gradual changes in frequency, phase, and symbol timing caused by drift and other changes in the plant. Figure 49 shows a method 4000 for tracking these changes. Steps 4010, executed by unit 153, Figure 22 or 155, Figure 23, track the downstream power level to compensate for slow gain changes in the plant. Step 4011 measures an average signal level by known methods from the output of FFT 170 or 180. If it is correct, exit 4012 takes no action. If the error is wrong but within a predetermined threshold of the correct value, block 4013 adjusts a coefficient of equalizer 172 to return the signal level to its nominal value. If the signal level changes more than a certain amount, or if it changes rapidly enough, the signal has probably been lost entirely. In that case, exit 4014 exits tracking mode; it may reenter the ranging procedure of Figures 47 and 48, or it may merely signal an error to the HDT.

0004374-070004  
T06020-42E70669

5

20

25

After the upstream transmitting modem 101 portion shown in Figure 24 has tuned to a subband in Figures 47 and 48, it and the upstream receiving modem 82 portion of Figure 26 must continue to track in amplitude, frequency, and timing. The use of multicarrier (MC) technology imposes some requirements which are not

obvious from experience with TDMA (time-division multiple access) or other bidirectional multipoint technologies, nor from point-to-multipoint (broadcast) MC networks and systems. In TDMA and similar systems, slight errors in frequency and timing, and larger errors in signal amplitude, can be compensated by tracking the signal at the receiver. In broadcast MC systems, all carriers are inherently synchronized at the head end, and can be tracked adequately at each receiver separately. In a bidirectional multipoint-to-point multicarrier plant, however, the head-end receiver sees all channels as though they had been generated by a single source, because the HDT decodes all channels in an entire 6MHz band as a single orthogonal waveform, with a single FFT converter. Even slight errors among the various ISUs in their 10-channel and 130-channel subbands can cause severe distortion and intersymbol interference when the FFT in modem 82 portion of Figure 26 converts the channels back into bit strings for multiple DS0 channels. The parameters to be controlled are frequency, symbol timing, and signal amplitude or power level. Frequency and timing can be tracked in a manner similar to steps 4020 and 4030 of Figures 47 and 48. Amplitude, however, has been found to be more critical than in previous systems.

Figure 50 depicts a method 4100 for tracking changes in the upstream channel signal amplitude. After Figures 47 and 48 have brought the ISU transmitting modem of Figure 24 on line, its amplitude must be balanced with that of all other ISUs in the system. Again, if different upstream subbands were received by different hardware, or in a TDMA fashion, where amplitude tracking could be particularized in frequency and/or time, a conventional AGC circuit could track amplitude variations adequately. In the embodiment described, power variations greater than about 0.25dB from one subband to another cause a significant amount of distortion and intersymbol interference. In a physical system of, say, a 20km radius, variations in the upstream signal level at the head end may vary by 20dB or more for different ISU locations, and may additionally vary significantly over time because of temperature differences, channel loading, component aging, and many

other factors. Conventional methods cannot achieve both the wide dynamic range and the high resolution required for an MC bidirectional multipoint-to-point system.

The steps in the left column of Figure 50 are performed within each ISU; the HDT performs the steps on the right. Step 4110 selects a number of payload channels for monitoring from table 4111. The channels must include one channel from each separate ISU, but need not include more than one. An MISU thus needs time in only one of its 130 payload channels, a very low overhead. A 10-channel HISU subband, however, may require time in more than one of its 10 payload channels, because multiple HISUs can share the same subband. Of course, a powered-down ISU, or one having no upstream payload channels provisioned to it, need not participate in blocks 1740, because it does not transmit upstream at all. (It is alternatively possible to employ IOC channels instead of payload channels for this purpose. Although requiring less overhead, such use is generally much more complex to implement.)

After ranging procedure 3900 has acquired a correct initial power level, step 4120 performs a scan every 30msec. for all the selected payload channels, as indicated by arrow 4121. Each ISU responds at 4130 by sending a message 4131 on its selected upstream channel when the scan reaches that channel. In step, the HDT measures the received power level from each ISU separately. If the signal amplitude is within a certain range of its previous value, then steps 4150 compensate for the variation at the HDT. Step 4151 smooths the errors over several scans, to prevent sudden jumps from a single glitch. Step 4152 then adjusts the coefficients of equalizer 214 in the upstream receiving portion, Figure 26, of modem 82, Figure 3, to compensate for the variation. This sequence compensates for small, slow variations at a high resolution; the equalizer steps are small and very linear.

Step 4140 may detect that the HDT equalizer is far from nominal, near the end of its range — say, 4dB up or down from nominal, for an equalizer having a  $\pm 5$ dB range of 0.25dB steps. This condition might occur for a large number of accumulated small errors, or it could be caused by a sudden, major change in the

system, such as a break in an optical fiber followed by an automatic switch to another fiber having a different length. Steps 4160 compensate for this condition at the transmitting (ISU) end, rather than at the receiving (HDT) end. The HDT then stops the channel scan at step 4161, and step 4162 sends an amplitude-error message 4163 over an IOC channel, addressed to the offending ISU. The message specifies the amount and direction of the compensation to be applied. The ISU applies this correction at step 4170 to vary the power output of its transmitter 200, Figure 24. Conventional transmitter power controls, such as a PIN attenuator diode 201 in power amplifier 200, are typically relatively coarse and nonlinear, but they do possess a wide range of adjustment. DAC 203 receives IOC messages to control attenuator 201. To allow the head-end equalizer to track the changing power more easily, step 4162 applies the correction over a long period of time, for example, 4-5sec/dB; but, if the monitored channel is the only channel connected to that ISU, the entire correction can be made in a single large step. Instead of controlling adjustment speed at the ISU, the head end may send individual timed IOC messages for multiple partial corrections; the downside is increased message traffic on the IOC channels.

ISUs may come online and be powered down at odd times. To prevent a feckless attempt at correction when an ISU is powered down, or its signal has been lost for some other reason, step 4140 further detects a condition of substantially zero power received from the ISU. In that event, step 4180 sets the head-end equalizer to its default value and keeps it there.

Thus, power-leveling blocks 1740 take advantage of the characteristics of the system to adjust both ends in a way which achieves both high resolution and wide dynamic range. The digital control available in the head-end equalizer provides precision and linearity in tracking slow changes, and the analog control at the ISU provides a wide range, and still allows the head end to track out inaccuracies caused by its coarse and nonlinear nature.

Call processing in the communication system 10 entails the manner in which

a subscriber is allocated channels of the system for telephony transport from the HDT 12 to the ISUs 100. The present communication system in accordance with the present invention is capable of supporting both call processing techniques not involving concentration, for example, TR-8 services, and those involving concentration, such as TR-303 services. Concentration occurs when there are more ISU terminations requiring service than there are channels to service such ISUs. For example, there may be 1,000 customer line terminations for the system, with only 240 payload channels which can be allocated to provide service to such customers.

Where no concentration is required, such as for TR-8 operation, channels within the 6 MHz spectrum are statically allocated. Therefore, only reallocation of channels shall be discussed further below with regard to channel monitoring.

On the other hand, for dynamically allocated channels to provide concentration, such as for providing TR-303 services, the HDT 12 supports on-demand allocation of channels for transport of telephony data over the HFC distribution network 11. Such dynamic allocation of channels is accomplished utilizing the IOC channels for communication between the HDT 12 and the ISUs 100. Channels are dynamically allocated for calls being received by a customer at an ISU 100, or for calls originated by a customer at an ISU 100. The CXMU 56 of HDT 12, as previously discussed, implements IOC channels which carry the call processing information between the HDT 12 and the ISUs 100. In particular, the following call processing messages exist on the IOC channels. They include at least a line seizure or off-hook message from the ISU to the HDT; line idle or on-hook message from the ISU to the HDT; enable and disable line idle detection messages between the HDT and the ISUs.

Call processing in the communication system 10 entails the manner in which a subscriber is allocated channels of the system for telephony transport from the HDT 12 to the ISUs 100. The present communication system in accordance with the present invention is capable of supporting both call processing techniques not involving concentration, for example, TR-8 services, and those involving



5

10

15

25

When a call is originated by a subscriber at the ISU side, each ISU 100 is responsible for monitoring the channel units for line seizure. When line seizure is

detected, the ISU 100 must communicate this change along with the PIN address code for the originating line to the CXMU 56 of the HDT 12 using the upstream IOC operation channel. Once the CXMU 56 correctly receives the line seizure message, the CXMU 56 forwards this indication to the CTSU 54 which, in turn, provides the necessary information to the switching network to set up the call. The CTSU 54 checks the availability of channels and allocates a channel for the call originated at the ISU 100. Once a channel is identified for completing the call from the ISU 100, the CXMU 56 allocates the channel over the downstream IOC channel to the ISU 100 requesting line seizure. When a subscriber returns on hook, an appropriate line idle message is sent upstream to the HDT 12 which provides such information to the CTSU 54 such that the channel can then be allocated again to support TR-303 services.

Idle channel detection can further be accomplished in the modem utilizing another technique. After a subscriber at the ISU 100 has terminated use of a data payload channel, the MCC modem 82 can make a determination that the previously allocated channel is idle. Idle detection may be performed by utilizing the equalization process by equalizer 214 (Figure 26) which examines the results of the FFT which outputs a complex (I and Q component) symbol value. An error is calculated, as previously described herein with respect to equalization, which is used to update the equalizer coefficients. Typically, when the equalizer has acquired the signal and valid data is being detected, the error will be small. In the event that the signal is terminated, the error signal will increase, and this can be monitored by signal to noise monitor 305 to determine the termination of the payload data channel used or channel idle status. This information can then be utilized for allocating idle channels when such operation of the system supports concentration.

The equalization process can also be utilized to determine whether an unallocated or allocated channel is being corrupted by ingress as shall be explained in further detail below with respect to channel monitoring.

The telephony transport system may provide for channel protection from

ingress in several manners. Narrowband ingress is a narrowband signal that is coupled into the transmission from an external source. The ingress signal which is located within the OFDM waveform can potentially take the entire band offline. An ingress signal is (most likely) not orthogonal to the OFDM carriers, and under worst case conditions can inject interference into every OFDM carrier signal at a sufficient level to corrupt almost every DS0+ to an extent that performance is degraded below a minimum bit error rate.

One method provides a digitally tunable notch filter which includes an interference sensing algorithm for identifying the ingress location on the frequency band. Once located, the filtering is updated to provide an arbitrary filter response to notch the ingress from the OFDM waveform. The filter would not be part of the basic modem operation but requires the identification of channels that are degraded in order to "tune" them out. The amount of channels lost as a result of the filtering would be determined in response to the bit error rate characteristics in a frequency region to determine how many channels the ingress actually corrupted.

Another approach as previously discussed with respect to the ingress filter and FFT 112 of the MCC upstream receiver architecture of Figure 26 is the polyphase filter structure. The cost and power associated with the filter are absorbed at the HDT 12, while supplying sufficient ingress protection for the system. Thus, power consumption at the ISUs 100 is not increased. The preferred filter structure involves two staggered polyphase filters as previously discussed with respect to Figures 31 and 32 although use of one filter is clearly contemplated with loss of some channels. The filter/transform pair combines the filter and demodulation process into a single step. Some of the features provided by polyphase filtering include the ability to protect the receive band against narrowband ingress and allow for scalable bandwidth usage in the upstream transmission. With these approaches, if ingress renders some channels unusable, the HDT 12 can command the ISUs to transmit upstream on a different carrier frequency to avoid such ingress.

The above approaches for ingress protection, including at least the use of

digital tunable notch filters and polyphase filters, are equally applicable to point to point systems utilizing multicarrier transport. For example, a single MISU transporting to a single HDT may use such techniques. In addition, uni-directional multi-point to point transport may also utilize such techniques for ingress protection.

5 In addition, channel monitoring and allocation or reallocation based thereon may also be used to avoid ingress. External variables can adversely affect the quality of a given channel. These variables are numerous, and can range from electro-magnet interference to a physical break in an optical fiber. A physical break in an optical fiber severs the communication link and cannot be avoided by  
10 switching channels, however, a channel which is electrically interfered with can be avoided until the interference is gone. After the interference is gone the channel could be used again.

Referring to Figure 40, a channel monitoring method is used to detect and avoid use of corrupted channels. A channel monitor 296 is used to receive events  
15 from board support software 298 and update a channel quality table 300 in a local database. The monitor 296 also sends messages to a fault isolator 302 and to channel allocator 304 for allocation or reallocation. The basic input to the channel monitor is parity errors which are available from hardware per the DS0+ channels; the DS0+ channels being 10-bit channels with one of the bits having a parity or data  
20 integrity bit inserted in the channel as previously discussed. The parity error information on a particular channel is used as raw data which is sampled and integrated over time to arrive at a quality status for that channel.

Parity errors are integrated using two time frames for each of the different service types including POTS, ISDN, DDS, and DS1, to determine channel status.  
25 The first integration routine is based on a short integration time of one second for all service types. The second routine, long integration, is service dependent, as bit error rate requirements for various services require differing integration times and monitoring periods as seen in Table 4. These two methods are described below.

Referring to Figures 41, 42, and 43, the basic short integration operation is

Subaz

described. When a parity error of a channel is detected by the CXMU 56, a parity interrupt is disabled by setting the interrupt priority level above that of the parity interrupt (Figure 41). If a modem alarm is received which indicates a received signal failure, parity errors will be ignored until the failure condition ends. Thus, some failure conditions will supersede parity error monitoring. Such alarm conditions may include loss of signal, modem failure, and loss of synchronization. If a modem alarm is not active, a parity count table is updated and an error timer event as shown in Figure 42 is enabled.

When the error timer event is enabled, the channel monitor 296 enters a mode wherein parity error registers of the CXMU 56 are read every 10 milliseconds and error counts are summarized after a one second monitoring period. Generally, the error counts are used to update the channel quality database and determine which (if any) channels require re-allocation. The channel quality table 300 of the database contains an ongoing record of each channel. The table organizes the history of the channels in categories such as: current ISU assigned to the channel, start of monitoring, end of monitoring, total error, errors in last day, in last week and in last 30 days, number of seconds since last error, severe errors in last day, in last week and in last 30 days, and current service type, such as ISDN, assigned to the channel.

As indicated in Figure 41, after the parity interrupt is disabled and no active alarm exists, the parity counts are updated and the timer event is enabled. The timer event (Figure 42), as indicated above, includes a one second loop where the errors are monitored. As shown in Figure 42, if the one second loop has not elapsed, the error counts are continued to be updated. When the second has elapsed, the errors are summarized. If the summarized errors over the one second period exceed an allowed amount indicating that an allocated channel is corrupted or bad, as described below, channel allocator 304 is notified and ISU transmission is reallocated to a different channel. As shown in Figure 43, when the reallocation has been completed, the interrupt priority is lowered below parity so that channel monitoring continues and the channel quality database is updated concerning the

5

10

15

20

25

Sub  
as

00001374 030901  
105020 120000

**Abstract**

10

15

20

25

8 (data) bits of error. Second, each monitoring parity error represents 80 frames of data (10 ms/125  $\mu$ s). Since the parity error is latched, all errors will be detected, but multiple errors will be detected as one error.

The bit error rate (BER) used as a basis for determining when to reallocate a channel has been chosen as  $10^{-3}$ . Therefore, the acceptable number of parity errors in a one second interval that do not exceed  $10^{-3}$  must be determined. To establish the acceptable parity errors, the probable number of frame errors represented by each observed (monitored) parity error must be predicted. Given the number of monitored parity errors, the probable number of frame errors per monitored parity error, and the number of bit errors represented by a frame (parity) error, a probable bit error rate can be derived.

A statistical technique is used and the following assumptions are made:

1. Errors have a Poisson distribution, and
2. If the number of monitored parity errors is small ( $< 10$ ) with respect to the total number of "samples" (100), the monitored parity error rate (PER) reflects the mean frame error rate (FER).

Since a monitored parity error (MPE) represents 80 frames, assumption 2 implies that the number of frame errors (FEs) "behind" each parity error is equal to 80 PER. That is, for 100 parity samples at 10 ms per sample, the mean number of frame errors per parity error is equal to 0.8 times the count of MPEs in one second. For example, if 3 MPEs are observed in a one second period, the mean number of FEs for each MPE is 2.4. Multiplying the desired bit error rate times the sample size and dividing by the bit errors per frame error yields the equivalent number of frame errors in the sample. The number of FEs is also equal to the product of the number of MPEs and the number of FEs per MPE. Given the desired BER, a solution set for the following equation can be determined.

$$(MPE \frac{FE}{MPE}) = 0.8MPE$$

The Poisson distribution, as follows, is used to compute the probability of a given number of FEs represented by a MPE ( $\chi$ ), and assumption 2, above, is used to arrive at the mean number of FEs per MPE ( $\mu$ ).

$$P(x) = \frac{e^{-\mu} \mu^x}{x!}$$

Since the desired bit error rate is a maximum, the Poisson equation is applied successively with values for  $\chi$  of 0 up to the maximum number. The sum of these probabilities is the probability that no more than  $\chi$  frame errors occurred for each monitored parity error.

The results for a bit error rate of  $10^{-3}$  and bit errors per frame error of 1 and 8 are shown in Table 3.

Table 3: Bit Error Rate Probability

Bit Errors per Frame Error	Monitored Parity Errors	Maximum Frame Errors/Monitored Parity Error ( $\chi$ )	Average Frame Errors/Monitored Parity Error ( $\mu$ )	Probability of BER $< 10^{-3}$
8	2	4	1.6	98%
	3	3	2.4	78%
	4	2	3.2	38%
1	8	8	6.4	80%
	9	7	7.2	56%
	10	7	8.0	45%

Using this technique, a value of 4 monitored parity errors detected during a one second integration was determined as the threshold to reallocate service of an ISU to a new channel. This result is arrived at by assuming a worst case of 8 bit errors per frame error, but a probability of only 38% that the bit error rate is better



**S**

10

15

20

25

5

10

15

20

25

mode to see if the channel has improved. In loop back mode, the CXMU 56 commands an initialized and activated ISU 100 to transmit a message on the channel long enough to perform short or long integration on the parity errors as appropriate. In the loop back mode, it can be determined whether the previously corrupted channel has improved over time and the channel quality database is updated accordingly. When not in the loop back mode, such channels can be powered down.

As described above, the channel quality database includes information to allow a reallocation or allocation to be made in such a manner that the channel used for allocation or reallocation is not corrupted. In addition, the information of the channel quality database can be utilized to rank the unallocated channels as for quality such that they can be allocated effectively. For example, a channel may be good enough for POTS and not good enough for ISDN. Another additional channel may be good enough for both. The additional channel may be held for ISDN transmission and not used for POTS. In addition, a particular standby channel of very good quality may be set aside such that when ingress is considerably high, one channel is always available to be switched to.

In addition, an estimate of signal to noise ratio can also be determined for both unallocated and allocated channels utilizing the equalizer 214 of the MCC modem 82 upstream receiver architecture as shown in Figure 26. As described earlier, the equalizer was previously utilized to determine whether a channel was idle such that it could be allocated. During operation of the equalizer, as previously described, an error is generated to update the equalizer coefficients. The magnitude of the error can be mapped into an estimate of signal to noise ratio (SNR) by signal to noise monitor 305 (Figure 26). Likewise, an unused channel should have no signal in the band. Therefore, by looking at the variance of the detected signal within the unused FFT bin, an estimate of signal to noise ratio can be determined. As the signal to noise ratio estimate is directly related to a probable bit error rate, such probable bit error rate can be utilized for channel monitoring in order to determine whether a bad or good channel exists.

Sub  
ay  
5

00001374-070901  
T06020-12E70600

Therefore, for reallocation for nonconcentration services such as TR-8 services, reallocation can be performed to unallocated channels with such unallocated channels monitored through the loopback mode or by SNR estimation by utilization of the equalizer. Likewise, allocation or reallocation for concentration services such as TR-303 services can be made to unallocated channels based upon the quality of such unallocated channels as determined by the SNR estimation by use of the equalizer.

With respect to channel allocation, a channel allocator routine for channel allocator 304 examines the channel quality database table to determine which DS0+ channels to allocate to an ISU 100 for a requested service. The channel allocator also checks the status of the ISU and channel units to verify in-service status and proper type for the requested service. The channel allocator attempts to maintain an optimal distribution of the bandwidth at the ISUs to permit flexibility for channel reallocation. Since it is preferred that ISUs 100, at least HISUs, be able to access only a portion of the RF band at any given time, the channel allocator must distribute channel usage over the ISUs so as to not overload any one section of bandwidth and avoid reallocating in-service channels to make room for additional channels.

The process used by the channel allocator 304 is to allocate equal numbers of each ISU type to each band of channels of the 6 MHz spectrum. If necessary, in use channels on an ISU can be moved to a new band, if the current ISU band is full and a new service is assigned to the ISU. Likewise, if a channel used by an ISU in one band gets corrupted, the ISU can be reallocated to a channel in another subband or band of channels. Remember that the distributed IOC channels continue to allow communication between the HDT 12 and the HISUs as an HISU always sees one of the IOC channels distributed throughout the spectrum. Generally, channels with the longest low-error rate history will be used first. In this way, channels which have been marked bad and subsequently reallocated for monitoring purposes will be used last, since their histories will be shorter than channels which have been operating in

a low error condition for a longer period.

### Second Embodiment of Telephony Transport System

A second embodiment of an OFDM telephony transport system, referring to  
 5 Figures 36-39 shall be described. The 6 MHz spectrum allocation is shown in  
 Figure 36. The 6 MHz bandwidth is divided into nine channel bands corresponding  
 to the nine individual modems 226 (Figure 37). It will be recognized by one skilled  
 in the art that less modems could be used by combining identical operations. Each  
 of the channel bands includes 32 channels modulated with a quadrature 32-ary  
 10 format (32-QAM) having five bits per symbol. A single channel is allocated to  
 support transfer of operations and control data (IOC control data) for  
 communication between an HDT 12 and ISUs 100. This channel uses BPSK  
 modulation.

The transport architecture shall first be described with respect to downstream  
 15 transmission and then with respect to upstream transmission. Referring to Figure  
 37, the MCC modem 82 architecture of the HDT 12 will be described. In the  
 downstream direction, serial telephony information and control data is applied from  
 the CXMC 80 through the serial interface 236. The serial data is demultiplexed by  
 demultiplexer 238 into parallel data streams. These data streams are submitted to a  
 20 bank of 32 channel modems 226 which perform symbol mapping and fast Fourier  
 transform (FFT) functions. The 32 channel modems output time domain samples  
 which pass through a set of mixers 240 that are driven by the synthesizer 230. The  
 mixers 240 create a set of frequency bands that are orthogonal, and each band is  
 then filtered through the filter/combiner 228. The aggregate output of the  
 25 filter/combiner 228 is then up-converted by synthesizer 242 and mixer 241 to the  
 final transmitter frequency. The signal is then filtered by filter 232, amplified by  
 amplifier 234, and filtered again by filter 232 to take off any noise content. The  
 signal is then coupled onto the HFC distribution network via telephony transmitter  
 14.

On the downstream end of the HFC distribution network 11, an ISU 100 includes a subscriber modem 258 as shown in Figure 38. The downstream signals are received from the ODN 18 through the coax leg 30, and are filtered by filter 260 which provides selectivity for the entire 6 MHz band. Then the signal is split into two parts. The first part provides control data and timing information to synchronize clocks for the system. The second part provides the telephony data. With the control data received separately from the telephony data, this is referred to as previously described above as an out of band ISU. The out of band control channel which is BPSK modulated is split off and mixed to baseband by mixer 262. The signal is then filtered by filter 263 and passed through an automatic gain control stage 264 and a Costas loop 266 where carrier phase is recovered. The signal that results is passed into a timing loop 268 so timing can be recovered for the entire modem. The IOC control data, which is a byproduct of the Costas loop, is passed into the 32 channel OFDM modem 224 of the ISU 100. The second part of the downstream OFDM waveform is mixed to base band by mixer 270 and associated synthesizer 272. The output of the mixer 270 is filtered by filter 273 and goes through a gain control stage 274 to prepare it for reception. It then passes into the 32 channel OFDM modem 224.

Referring to Figure 39, the IOC control data is hard limited through function block 276 and provided to microprocessor 225. The OFDM telephony data is passed through an analog to digital converter 278 and input to a first-in first-out buffer 280 where it is stored. When a sufficient amount of information is stored, it is accepted by the microprocessor 225 where the remainder of the demodulation process, including application of an FFT, takes place. The microprocessor 225 provides the received data to the rest of the system through the receive data and receive data clock interface. The fast Fourier transform (FFT) engine 282 is implemented off the microprocessor. However, one skilled in the art will recognize that the FFT 282 could be done by the microprocessor 225.

In the upstream direction, data enters the 32 channel OFDM modem 224

through the transmit data ports and is converted to symbols by the microprocessor 225. These symbols pass through the FFT engine 282, and the resulting time domain waveform, including guard samples, goes through a complex mixer 284. The complex mixer 284 mixes the waveform up in frequency and the signal is then passed through a random access memory digital to analog converter 286 (RAM DAC). The RAM DAC contains some RAM to store up samples before being applied to the analog portion of the ISU upstream transmitter (Figure 38). Referring to Figure 38, the OFDM output for upstream transport is filtered by filter 288. The waveform then passes through mixer 290 where it is mixed under control of synthesizer 291 up to the transmit frequency. The signal is then passed through a processor gain control 292 so that amplitude leveling can take place in the upstream path. The upstream signal is finally passed through a 6 MHz filter 294 as a final selectivity before upstream transmission on the coaxial leg 30 to the ODN 18.

In the upstream direction at the HDT 12 side, a signal received on the coax from the telephony receiver 16 is filtered by filter 244 and amplified by amplifier 246. The received signal, which is orthogonally frequency division multiplexed, is mixed to baseband by bank of mixers 248 and associated synthesizer 250. Each output of the mixers 248 is then filtered by baseband filter bank 252 and each output time domain waveform is sent then to a demodulator of the 32 channel OFDM modems 226. The signals pass through a FFT and the symbols are mapped back into bits. The bits are then multiplexed together by multiplexer 254 and applied to CXMC 56 through the other serial interface 256.

As shown in this embodiment, the ISU is an out of band ISU as utilization of separate receivers for control data and telephony data is indicative thereof as previously discussed. In addition, the separation of the spectrum into channel bands is further shown. Various other embodiments as contemplated by the accompanying claims of the transport system are possible by building on the embodiments described herein. In one embodiment, an IOC control channel for at least synchronization information transport, and the telephony service channels or paths

are provided into a single format. The IOC link between the HDT 12 and the ISUs 100 may be implemented as four BPSK modulated carriers operating at 16 kbps, yielding a data rate of 64 kbps in total. Each subscriber would implement a simple separate transceiver, like in the second embodiment, which continuously monitors the service channel assigned to it on the downstream link separately from the telephony channels. This transceiver would require a tuned oscillator to tune to the service IOC channel. Likewise, an IOC channel could be provided for channel bands of the 6 MHz bandwidth and the channel bands may include orthogonal carriers for telephony data and an IOC channel that is received separately from the reception of the orthogonal carriers.

In another embodiment, instead of 4 BPSK channels, a single 64 kbps IOC channel is provided. This single channel lies on the OFDM frequency structure, although the symbol rate is not compatible with the telephony symbol rate of OFDM framework. This single wide band signal requires a wider band receiver at the ISU 100 such that the IOC link between the HDT 12 and ISUs is always possible. With single channel support it is possible to use a fixed reference oscillator that does not have to tune across any part of the band in the subscriber units. However, unlike in the first embodiment where the IOC channels are distributed across the spectrum allowing for narrow band receivers, the power requirements for this embodiment would increase because of the use of the wide band receiver at the ISU 100.

In yet another embodiment, the IOC link may include two IOC channels in each of 32 OFDM channel groups. This increases the number of OFDM carriers to 34 from 32 in each group. Each channel group would consist of 34 OFDM channels and a channel band may contain 8 to 10 channels groups. This approach allows a narrow band receiver to be used to lock to the reference parameters provided by the HDT 12 to utilize an OFDM waveform, but adds the complexity of also having to provide the control or service information in the OFDM data path format. Because the subscriber could tune to any one of the channel groups, the information that is embedded in the extra carriers must also be tracked by the central office. Since the



system needs to support a timing acquisition requirement, this embodiment may also require that a synchronization signal be placed off the end of the OFDM waveform.

It is to be understood, however, that even though numerous characteristics of the present invention have been set forth in the foregoing description, together with details of the structure and function of the invention, the disclosure is illustrative and changes in matters of order, shape, size, and arrangement of the parts, and various properties of the operation may be made within the principles of the invention and to the full extent indicated by the broad general meaning of the terms in which the appended claims are expressed.

In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings that form a part hereof, and in which are shown by way of illustration specific embodiments in which this invention may be practiced. It is understood that other embodiments may be used and structural changes may be made without departing from the scope of the claimed invention.

Figure 70 shows one embodiment of this invention having an apparatus, generally indicated as FFT system 2100, which performs both forward and inverse FFT functions. The input/output signals to FFT system 2100 include some (or, in one preferred embodiment, all) of the following: real data-in 2111 (having an N-bit-wide data path; in one embodiment, this is 10 bits wide, and bi-directional so it can be both written to and read from), imaginary data-in 2112 (having an N-bit-wide data path; in one embodiment, this is 10 bits wide, and bi-directional so it can be both written to and read from), address in 2113 (having enough bits to specify addresses for each input value or input-value pair), control and clock lines 2114 that control writing (and reading) data from the input side, test signals 2115, size-select bits 2116, bit-growth-select bits for each of P stages including bit-growth-1 bits 2117 through bit-growth-P bits 2118, forward/inverse select signal 2119 which specifies whether to perform a forward or inverse transform, power-down command signal input 2109, real data-out 2121 (having an M-bit-wide data path, in one embodiment, this is 10 bits), imaginary data-out 2122 (having an M-bit-wide data

5

## 10

15

## 20

25

Figure 71 is a block diagram of modem 2400 according to the present invention which includes a FFT system 2100 configured to perform an IFFT in

5  
10  
15  
20  
25

In one embodiment, the functions of all three banks of RAMs are also "permuted" on the rising edge of this 8-KHz frame clock 2108, before each FFT

calculation starts. As used herein, this "permutation" changes the function of each bank of RAM without actually moving data: input RAM 2251 becomes conversion RAM 2252, conversion RAM 2252 becomes output RAM 2253, and output RAM 2253 becomes input RAM 2251. Note that, in this embodiment, no RAM data is moved when the functions are permuted. Figure 73 shows one embodiment of a physical implementation which provides the function of input RAM 2251, conversion RAM 2252, and output RAM 2253. Bank control block 2131 permutes the function of the physical RAM banks 2151, 2152 and 2153 at the rising edge of frame clock 2139. One of the functions of bank control state machine 2131 is to control the routing of data through RAM input multiplexers (MUXs) 2171, 2172, and 2173 and the routing of data through RAM output multiplexers (MUXs) 2181, 2182, and 2183. For example, at a first given state, bank control 2131 controls input-select block 2132 to input data into physical RAM bank 2151 (in one embodiment, physical RAM bank 2151 includes 1024 ten-bit real values and 1024 ten-bit imaginary values). Thus, in the first state, physical RAM bank 2151 acts as logical input RAM bank 2251. Once all of the desired first set of input values (up to 1024 values or pairs of values) have been inputted, the frame clock 2139 is driven to change the state of bank control 2131 (permuting the three RAM bank's respective functions) to a second state, in which bank control 2131 controls input-select block 2132 to input data into physical RAM bank 2152 and bank control 2131 controls conversion-select block 2133 to direct computation accesses for data into physical RAM bank 2151. (The five low-order bits of values used for computation are always provided from physical RAM bank 2154.) Thus, in the second state, physical RAM bank 2152 acts as logical input RAM bank 2251 and physical RAM bank 2151 acts as logical conversion RAM bank 2252. Once all of the desired second set of input values (up to 1024 values or pairs of values) have been inputted, the frame clock 2139 is driven to change the state of bank control 2131 (permuting the three RAM bank's respective functions again) to a third state, in which bank control 2131 controls input-select block 2132 to input data into physical RAM bank

2153, bank control 2131 controls conversion-select block 2133 to direct computation accesses for data into physical RAM bank 2152, and bank control 2131 controls output-select block 2134 to direct output requests for accesses for data from physical RAM bank 2151. Thus, in the third state, physical RAM bank 2153 acts as logical input RAM bank 2251, physical RAM bank 2152 acts as logical conversion RAM bank 2252, and physical RAM bank 2151 acts as logical output RAM bank 2253. Once all of the desired third set of input values (up to 1024 values or pairs of values) have been inputted – and the first set of converted output values (up to 1024 values or pairs of values) have been outputted – then the frame clock 2139 is driven to change the state of bank control 2131 (permuting the three RAM bank's respective functions again) back to the first state, in which bank control 2131 controls input-select block 2132 to input data into physical RAM bank 2151, bank control 2131 controls conversion-select block 2133 to direct computation accesses for data into physical RAM bank 2152, and bank control 2131 controls output-select block 2134 to direct output requests for accesses for data from physical RAM bank 2153. Thus, in the first state, physical RAM bank 2151 again acts as logical input RAM bank 2251, physical RAM bank 2153 acts as logical conversion RAM bank 2252, and physical RAM bank 2152 acts as logical output RAM bank 2253.

At the beginning of each transform (i.e., FFT/IFFT (fast Fourier transform or inverse fast Fourier transform)) process, the functions of the RAM bank memories (the mapping of RAMs 2151, 2152, and 2153 to RAMs 2251, 2252, and 2253) are permuted. The conversion RAM 2252 becomes the output RAM 2253, the input RAM 2251 becomes the conversion RAM 2252, and the output RAM 2253 becomes the input RAM 2251. Each RAM (2251, 2252, and 2253) has its own independent control and clock signals (2114, 2128, and 2124, respectively). While the current transform is being calculated in conversion RAM 2252, the results of the previously calculated FFT/IFFT may be read from the output RAM 2253 and the data for the next FFT calculation may be asynchronously and simultaneously written to the input RAM 2251. In one embodiment, access to the input RAM 2251 and output RAM

2253 is restricted for 3 clock cycles (one prior to and two after) relative to the rising edge of the frame clock. This allows the RAM permutation to proceed safely without undesirable data loss.

5      Forward Fast Fourier Transform (FFT) and  
       Inverse Fast Fourier Transform (IFFT)

10      In one embodiment, FFT system 2100 is implemented on a single integrated circuit (IC) that performs both a forward FFT and an inverse FFT ("IFFT", also called a reverse FFT). A pin, forward/inverse signal 2119, selects between the two types of transform. The inverse FFT uses the identical calculation sequence as the forward transform, but the complex values (i.e., real and imaginary) of the twiddle factors and butterfly coefficients are conjugated relative to these values used for the FFT. A forward FFT is defined to convert a time-domain signal into frequency-domain signals and, in one embodiment, is used in the transmitter 2401 of modem 2400 (see Figure 71). An inverse FFT is defined to convert frequency-domain signals to time-domain signals and is used in the receiver 2402 of the modem 2400.

15      In one embodiment, the underlying structure of FFT system 2100 supports five radix-4 butterflies (i.e., butterfly operations) usable to perform the 1024-point transforms (FFT and IFFT). The butterflies and stages are reduced for lower-order transforms. (In one embodiment, every other twiddle factor of the 1024-point twiddle-factor lookup table 2610 is used when calculating the 512-point transforms by forcing to zero the low-order address bit to the twiddle-factor lookup table 2610. Every fourth twiddle factor of the 1024-point twiddle-factor lookup table 2610 is used when calculating the 256-point transforms by forcing to zero the two low-order address bits to the twiddle-factor lookup table 2610.) In one embodiment, all FFT transforms are calculated by using radix-4 butterflies, except for the last stage of the 512-point, 128-point and 32-point transforms, which use the radix-4 structure to perform a radix-2 butterfly. (Obvious extensions are made if other transforms are used). Thus, the 1024-point transforms use five stages of radix-4 butterflies, the

20

25

512-point transforms use four stages of radix-4 butterflies followed by one stage of a radix-2 butterfly, etc.

#### Scaling of FFT output

5 In one embodiment, scaling is controlled by ten external pins (bit-growth signals 2117 through 2118) on the ASIC 2101. Two pins are used for each of the five passes required for the 1024- and 512-point transform. The use of two pins at each stage allows a scaling factor (e.g., a shift right of each intermediate result value) of from 0 bits to 3 bits which exceeds the nominal bit growth of 1 bit  
10 observed in the inventor's analysis. In one embodiment, the binary number represented by each pair of pins indicates the number of binary places that the results of each calculation are to be shifted right (divided) by before they are placed back into conversion RAM 2252. Note that in one such embodiment scaling the result after each butterfly operation requires at least some working registers to  
15 maintain bits greater than the MSB of the 15-bit values (each for real and imaginary) in the internal conversion RAM 2252, and the scaling to be applied prior to storage in the conversion RAM 2252.

#### Round-off/truncation, saturation, and scaling

20 In one embodiment, each input and output number is represented by 10 bits (i.e., 10 bits for each real portion, plus 10 bits for each imaginary portion of a complex pair of numbers). Numbers are represented as two's-complement fractional arithmetic with the binary point immediately to the right of the sign bit. Fractional arithmetic helps prevent multiplies from causing overflow. The result  
25 portion of a double-precision multiply which is used is the upper fifteen bits.

Number growth is managed by the 10 scaling pins (bit growth signals 2117-2118) mentioned previously. In addition, the ASIC 2101 uses saturation logic that prevents number roll over. That is to say, that if the result of an addition or subtraction exceeds the maximum value that can be represented in 15 bits, the

106020-12370600

result is replaced by a value representing the maximum possible integer. Likewise, if the result of an addition or subtraction is less than the minimum value that can be represented in 15 bits, the result is replaced by a value representing the minimum possible integer.

5 In one embodiment, ASIC 2101 has the following registers for real and complex data: 10-bit input RAM 2251 and 10-bit output RAM 2253, and a 15 bit conversion RAM 2252. The input values are presented as a two's complement fractional binary value (2.bbbbbbbb). In the internal conversion RAM, the five extra bits are appended at the bottom (i.e., low-order position) of the input word.

10 The resulting format is an extended 2's-complement value in the internal conversion RAM 2252 ( e.g., s.bbbbbbbbxxxx, where s is sign, b's are significant input bits, and x's are extra bits to maintain precision). Note that in one embodiment any extension bits in the working registers extend the sign bit to achieve the desired 2's-complement results.

15 In one embodiment, at the output of each transform pass, the result is shifted down as specified by the signals provided into the scaling pins 2117-2118 for that pass. The output RAM 2253 aligns precisely with the upper 10 bits of the conversion RAM 2252, in a manner similar to the input RAM 2251 (s.bbbbbbbbxxxx), each output value (after the final growth factor is applied prior

20 to storage in the internal conversion RAM 2252 is rounded to produce the output results (s.bbbbbbbbr).

#### Power Down

25 In one embodiment, the ASIC has a single pin, power-down signal 2109, to control power down that gates off the internal clock in the ASIC, and places the input RAM 2251 and output RAM 2253 into a low power state. In one embodiment, access attempts to those RAMS 2251 and 2253 while in the low power state will be unsuccessful.



### Execution Time

The maximum FFT execution frequency requires that a 1024-point transform be completed in 125 microseconds based on a 32 MHz clock.

An approximately 8-KHz frame clock 2108 is provided to the ASIC 2101 to signify the beginning of a transform cycle. Prior to the rising edge of this clock, the previous FFT should be complete. Table 5 below summarizes the clock frequency to perform a 1024-point and a 512-point FFT in 125 microseconds.

**Table 5: FFT Clock required to perform a transform in 125 microseconds**

Transform Size	Clock cycles per conversion	Clock Frequency for 125 $\mu$ Sec conversion rate (MHz)
1024	3845	30.8
512	1925	15.4
256	775	6.2
128	390	3.1
64	150	1.2
32	80	0.64

### FFT Functional Blocks

The ASIC is partitioned into the following major functional blocks.

1. RAM banks 2151, 2152, 2153
2. Sequencer 2640
3. Dual radix (2, 4) core 2600, including multipliers 2620 through 2627
4. Twiddle-factor lookup table 2610

The algorithm used is a decimation in time (DIT) FFT. The algorithm used is an in place algorithm which means that the results of each butterfly are put back into the same locations that the operands came from. The algorithm assumes digit

reversed input order and normal output order. Since the three banks of RAMs are independent, the reordering of the input data is done internally to the ASIC. This reordering is transparent to the user. Thus the user writes the data into the input RAM bank in normal order and read the data from the output RAM bank in normal order. If the user requires some other ordering, this may be accomplished by simply permuting the address line ordering.

The order of the memory for all transform sizes are given in the section titled "Sequencer."

#### Memory Banks

The ASIC 2101 includes three banks of RAM that are individually addressable. The RAM banks 2251, 2253, and 2252 are used for input data, output data and conversion data, respectively. The three banks are used in order to obtain real-time execution of the FFT and to allow independent clock rates for the input and output of data in order to accommodate clocking needs of those devices connected to it. In one embodiment, the input RAM bank 2251 and output RAM bank 2252 are each organized as 1024 twenty-bit words. The lower 10 bits of each 20-bit word are used for real data, and the upper 10 bits of each 20-bit word are used for imaginary data. The conversion RAM bank 2252 is organized as 1024 thirty-bit words. In one embodiment, the lower 15 bits of each thirty-bit word are used for real data and the upper 15 bits for imaginary data. The connectivities of the three RAM banks of memory are controlled by a state machine which is advanced by the 8-KHz frame signal 2108. The state machine has three states which are defined in Table 6 as follows:

**Table 6: Typical processing states of memory banks.**

State	BANK A	BANK B	BANK C
0	Input	Convert	Output
1	Convert	Output	Input
2	Output	Input	Convert

In one embodiment, the RAM banks do not have a reset. For input and output, they are accessed at a maximum of 10.24 MHz rate. In one embodiment, a dead time of one clock cycle before the rising edge of the 8-KHz clock 2108 and two clock cycles after is required to assure the safety of the RAM bank switching. The ASIC 2101 clocks the data in on the rising edge of the input clock and clocks data out on the rising edge of the output clock.

Actual operation of the conversion RAM 2252 is a bit more complicated than is implied above. Actually only the top 10 bits of the conversion RAM 2252 participates in the bank switching between RAMs 2151, 2152, and 2153. The lower 5 bits 2154 are dedicated the conversion RAM 2252. Since all numbers are MSB aligned, no shifting needs to be done on input or output. However, the last stage of the FFT calculation rounds the results to ten bits. This eliminates a bias that would result had the results simply been truncated. The rounding is accomplished by adding a one to the eleventh bit position and then truncating this result.

#### Sequencer 2640

The sequencer 2640 manages the processing of the FFT system 2100. The sequencer 2640 controls the generation of addresses for the conversion RAM bank 2252 and the twiddle factor ROMs 2610 through address generation blocks 2642 and 2641 respectively. Sequencer 2640 also configures the calculation commands for the radix butterfly calculator 2630. In addition, sequencer 2640 monitors the calculations for overflow. If at any time during the course of FFT calculation, an overflow or underflow is detected, then a flag is set indicating that the results of the

FFT are suspect. This overflow flag is passed along with the output data block when the RAM banks are switched. Thus the flag pin 2125 indicates that the output data block presently being read out of the output RAM bank 2253 may not be accurate. This flag 2125 may aid the system designer in providing an indicator for AGC (automatic gain control).

#### Dual radix (2, 4) core 2600

The dual radix core 2600 is the arithmetic element of the ASIC 2101. In one preferred embodiment, it includes eight 16-bit-by-16-bit multipliers (2620 through 2627) and thirty-two multiplexed adder-subtractor-accumulators 2633. In another embodiment, it includes twelve 16-bit-by-16-bit multipliers and eight multibranched adder/subtractor/accumulators.

Figure 74 shows one embodiment of a dual radix core 2600. In the embodiment shown in Figure 74, conversion RAM bank 2252 is shown as part of the dual radix core. In other embodiments, conversion RAM bank 2252 is a separate functional unit, not considered part of the dual radix core 2600. In Figure 74, the data fetched from conversion RAM bank 2252 is fetched into holding latches 2612, with both the real (i.e., X3R, X2R, X1R, and X0R) and imaginary parts (i.e., X3I, X2I, X1I, and X0I) of four points being fetched in parallel, substantially simultaneously. As shown in Figure 74, the real and imaginary data fetched from conversion RAM bank 2252 is addressed by address generation circuit 2642 which provides four separate addresses to conversion RAM bank 2252. The real and imaginary part of a single point both use the same address. The four separate addresses allow different sets of points to be fetched simultaneously during various stages of the transform operation.

As further shown in Figure 74, the data fetched from conversion RAM bank 2252 includes real-and-imaginary pairs denoted X3, X2, X1 and X0. The data point pairs being denoted are X3 which includes X3R and X3I, X2 which includes X2R and X2I, X1 which includes X1R and X1I, and X0 which includes X0R and X0I. In

one embodiment, the data fetched from conversion RAM bank 2252 is held for two successive multiplier clock cycles in holding latches 2612, so that the same value gets multiplied in each of two multiply cycles. The corresponding twiddle factors are fetched on every clock cycle from twiddle factor look-up table 2610. For

5 example in one embodiment, on even clock cycles the real components for four twiddle factors are fetched in parallel denoted W0R, W1R, W2R, and W3R. These four real twiddle factors are fed to multipliers 2620 and 2621, 2622 and 2623, 2624 and 2625, and 2626 and 2627, respectively as shown, and multiplied by the real and

10 imaginary components of all four data value pairs fetched from conversion RAM bank 2252. The products of these eight multipliers 2620-2627 are fed through routing logic 2634 into all thirty-two adder/subtractor accumulators 2633 of row-column array 2632. On the successive odd multiplier clock cycle, four imaginary

15 twiddle factors are then fetched, shown as W0I, W1I, W2I and W3I. As shown, imaginary twiddle factor W0I is fed to both multiplier 2620 and multiplier 2621. Similarly, imaginary twiddle factor W1I is fed to both multiplier 2622 and multiplier 2623. Imaginary twiddle factor W2I is fed to multiplier 2624 and 2625

20 and imaginary twiddle factor W3I is fed to multiplier 2626 and 2627. For the second multiplier clock, these eight multipliers are again fed with the same complex point data still held in holding latches 2612 as was used in the first clock multiplier cycle. Again after the second clock multiplier, the eight products are fed through routing logic 2634 and accumulated in adder/subtractor accumulators 2633. Note

25 that each twiddle factor value in twiddle factor lookup table 2610 is a complex number having a real and imaginary part (in this embodiment, the real and imaginary parts are stored in successive locations which are fetched on successive clock cycles, thus presenting 4 real values followed by four imaginary values) and

each point value in conversion RAM bank 2252 is also a complex number having a real part and an imaginary part (in this embodiment, the real and imaginary parts are stored in side-by-side locations which are fetched on only every other clock cycle, thus presenting 4 real values and four imaginary values on every other clock cycle).

FIG. 20

These elements are configured to perform a radix-4 butterfly calculation. The radix-4 butterfly calculations needed to provide a 1024-point FFT or IFFT and the selection of addresses and twiddle coefficients for the, e.g., 1024 points used for each of the five passes needed for a complete 1024-point transform are easily calculated and are well known in the art (see, for example, chapters 6 & 10 of *Theory and Application of Digital Signal Processing*, by Lawrence R. Rabiner and Bernard Gold, Prentice-Hall Inc, published in 1975; and page 612 et seq. of *Discrete-Time Signal Processing*, by Alan V. Oppenheim and Ronald W. Schaffer, Prentice-Hall Inc, published in 1989). Although a radix-2 transform is also required, a separate radix-2 calculator is not required because the radix-2 operation is a subset of the radix-4 operation. Only the operands actually needed for the radix-2 operation are actually loaded. The remainder are set to zero. All arithmetic is performed in such a manner that, if an overflow or underflow should occur, then the results saturate. This prevents roll-over from contaminating the results. The function performed by the dual-radix core 2600 for a radix-4 butterfly operation supporting the forward transform is the following:

$$\begin{aligned} A' &= A + BW^{1k} + CW^{2k} + DW^{3k} \\ B' &= A - jBW^{1k} - CW^{2k} + jDW^{3k} \\ C' &= A - BW^{1k} + CW^{2k} - DW^{3k} \\ D' &= A + jBW^{1k} - CW^{2k} - jDW^{3k} \end{aligned}$$

where A, B, C and D are the four input points, A', B', C' and D' are the four output points, and  $W^{1k}$ ,  $W^{2k}$  and  $W^{3k}$  are the twiddle coefficients.

The function performed in a reverse transform is simply the complex conjugate of the above set of equations.

$$\begin{aligned} A' &= A + BW^{1k} + CW^{2k} + DW^{3k} \\ B' &= A + jBW^{1k} - CW^{2k} - jDW^{3k} \\ C' &= A - BW^{1k} + CW^{2k} - DW^{3k} \end{aligned}$$

$$D' = A - jBW^{14} - CW^{24} + jDW^{24}$$

Note that the twiddle-factor  $W$  which is used for each  $A$  on the right side of the above equations is one (i.e., the complex number  $1 + j0$ ).

The multiplication of two complex numbers, each having a real part and an imaginary part, for example results in the following equation:

$$\begin{aligned} X0 \times W0 &= (X0R + jX0I) \times (W0R + jW0I) = \\ &= (X0R \times W0R - X0I \times W0I) + j(X0I \times W0R + X0R \times W0I) = \\ &+ X0R \times W0R + jX0I \times W0R \text{ (e.g., the first multiplier cycle in multipliers 2620-2621)} \\ &- X0I \times W0I + jX0R \times W0I \text{ (e.g., the first multiplier cycle in multipliers 2620-2621)}. \end{aligned}$$

Thus four multiplier operations are needed for each complex multiply operation.

In order to speed the transform function, the factors for the parallel multipliers are fetched in parallel under the control of control and clocking sequencer 2640. Routing logic 2634 routes the products of the multipliers 2620-2627 to the thirty-two adder-subtractor-accumulators 2633. In one embodiment, eight multiplier cycles, C0 through C7, are used to generate four radix-4 butterfly operations, resulting in sixteen complex output values. For discussion purposes, row-column array 2632 is shown having four rows (A, B, C, and D) and four columns (W, X, Y, and Z) of complex value pairs. The real value and the imaginary value of each of these sixteen complex value pairs has its own associated adder-subtractor-accumulator 2633, for a total of thirty-two adder-subtractor-accumulators 2633, as shown in Figure 74 and Figure 91. In one embodiment, scaling-factor shift logic 2644 (under the control of bit-growth selector 2643) is provided in the path between adder-subtractor-accumulators 2633 and conversion RAM bank 2252. The scaling-factor shift logic 2644 provides a right-shift function of 0 bits, 1 bit, 2 bits or three bits (divide by 1, two, four or eight respectively) on each output data value as it is being returned to conversion RAM 2252. Bit-growth pins 2117 through 2118, which control the divide function for each of the passes are coupled to bit-growth selector 2643 under control of sequencer 2640.

Figures 75-82 are a table 2800 showing the order of calculations for a

§

10

15

20

25



be the WR value loaded in cycle C0) the output (called WI) of multiplier 2621, and the adder-subtractor-accumulator 2633 for the imaginary portion of the DZI point in row-column array 2632 (called the DZI accumulator) gets loaded with the sum of its previous value (called DZI, this value happens to be the ZI value loaded in cycle C0) and output (called ZR) of multiplier 2626.

Similarly, Figures 77 through 82 show the operations which take place at multiplier clocks C2 through C7, respectively.

Since each complex-multiply operation takes a total of four multiplier operations, and two multipliers (e.g., the pair 2620 and 2621) are used, two multiplier cycles are needed for each complex-multiply operation. In a 1024-point transform (i.e., either an FFT or an IFFT), four of the five passes involve sets of four points wherein all four points are in a single sub-bank (e.g., 2290), and therefore must be fetched on four successive even-clocks. Each of these four passes takes eight clocks, called C0 through C7. These four passes are each called "normal butterfly." Table 2800 shows the order of calculation for all of the suboperations for one embodiment of a normal butterfly (calculating four radix-4 butterfly operations in eight multiplier clock cycles), where each of the four points for one radix-4 butterfly are in the same sub-bank (e.g., either sub-bank 2290 or 2291 or 2292 or 2293).

Figures 83-90 are a table 2810 showing the order of calculations for a "transposed butterfly sub-operation." The transposed butterfly sub-operation is used for one pass of each transform performed. The data points in conversion RAM 2252 are arranged within conversion RAM 2252 such that the four input points for one radix-4 operation are each located in different sub-banks if the points are successively addressed (e.g., addresses 0, 1, 2, and 3 are each in different sub-banks, e.g., sub-banks 2290, 2291, 2292, and 2293 respectively). The transposed butterfly passes for this one set of points (those whose addresses-mod-4 are equal) are handled by the equations shown in the table of Figures 83-90. Figure 83 shows the operations at each of the thirty-two adder-subtractor-accumulators 2633 at a

multiplier clock cycle command denoted C0 (note that only eight adder-subtractor-accumulators 2633 are affected, the other twenty-four do nothing). For example, at C0, the adder-subtractor-accumulator 2633 for the real portion of the AW point in row-column array 2632 (called the AWR accumulator) gets loaded with the result of a four-way addition of the outputs (called  $WR + XR + YR + ZR$ , these are the real-times-real portions) of multipliers 2620, 2622, 2624 and 2626, and the adder-subtractor-accumulator 2633 for the imaginary portion of the AZ point in row-column array 2632 (called the AZI accumulator) gets loaded with the sum/difference of outputs (called  $WI - XR - YI + ZR$ ) of multipliers 2621, 2622, 2625, and 2626, respectively. By performing load operations at clock C0 with no accumulation of the prior value (e.g., in AWR), the previous values of the accumulators do not need to be zeroed. Note that, since all four points for a single butterfly operation can be fetched simultaneously from conversion RAM 2252, and the results of the respective multiply operations must all be combined as they are formed, five-way mixed add/subtract operations are provided for by each adder-subtractor-accumulator 2633.

Figure 84 shows the operations at each of the thirty-two adder-subtractor-accumulators 2633 at a multiplier clock cycle command denoted C1. For example, at C1, the adder-subtractor-accumulator 2633 for the real portion of the AW point in row-column array 2632 (called the AWR accumulator) gets loaded with the result of a five-way subtraction/addition of the outputs (called  $AWR - (WI + XI + YI + ZI)$ , these are the imaginary-times-imaginary portions) of multipliers 2621, 2623, 2625 and 2627 and the prior contents of AWR. The adder-subtractor-accumulator 2633 for the imaginary portion of the AZ point in row-column array 2632 (called the AZI accumulator) gets loaded with the sum/difference of outputs (called  $AZI + (WR + XI - YR - ZI)$ ) of multipliers 2620, 2623, 2624, and 2627, respectively.

Similarly, Figures 85 through 90 show the four-way and five-way operations which take place at multiplier clocks C2 through C7, respectively. Table 2810 shows the order of calculation for one embodiment of a transposed butterfly

TOP SECRET - CONFIDENTIAL

(calculating four radix-4 butterfly operations in eight multiplier clock cycles), where each of the four points for one radix-4 butterfly are each in different sub-banks (e.g., one point in sub-bank 2290, one point in 2291, one point in 2292, and one point in 2293).

#### Twiddle-Factor Lookup Table 2610

In one embodiment, the twiddle-factor lookup table 2610 (also called a sine-cosine ROM lookup table) comprises 512 fifteen-bit words, wherein four words can be fetched in parallel. Each complex twiddle factor value pair is fetched sequentially, wherein the first 15-bit word represents the real part of the twiddle factor value pair and the second 15-bit word represents the imaginary part, and four values are fetched simultaneously (i.e., four real values, having 60 bits total, are fetched on an even clock – e.g., clocks C0, C2, C4 or C6 – and four imaginary values are fetched on the following odd clock – e.g., clocks C1, C3, C5 or C7). In another embodiment, the twiddle factor lookup table 2610 comprises of 256 thirty-bit words. The upper 15 bits represent the real part of the twiddle factor whereas the lower 15 bits represent the imaginary part. Although 1024 complex-value pairs are required in order to produce a 1024-point FFT or IFFT, the values are not unique, and the number of twiddle factors was reduced by a factor of four by making use of the simple trigonometric identities in mapping 360 degrees of twiddle factors to a 90-degree lookup table. In one embodiment, the twiddle-factor lookup table was designed to minimize DC offset caused by integer-based twiddle factors.

Figure 92 shows a more-detailed block diagram of an adder-subtractor-accumulator 2633. In one embodiment, multipliers 2620 through 2627 are each a 16-bit-by-16-bit multiplier. In one embodiment, only the upper-order 16 bits of the resultant product are passed by MUX 2834. (In one embodiment, MUX 2834 is part of router logic 2634.) Adder-subtractor 2833 performs a five-way addition/subtraction as defined in Figures 83-90 and the two-way addition/subtraction as defined in Figures 75-82, under the control of sequencer 2640. In one embodiment, accumulator 2835 maintains enough bits above the

binary point to accomodate overflow bits and to provide an indication of overflow which does not get lost as further addition/subtractions are performed on the accumulating data (in other embodiments, one, two, or three such bits are maintained).

5

#### Input and Output Timing

Below is the detailed timing for one embodiment of the input and output RAMs.

10

**Table 7: Read Cycle for output RAM bank 2253**

NOTE: The RAM clocks for the input and output banks are limited to 10.24 MHz.

Symbol	Parameters	Condition	Nom
Tcc	Clock Cycle Time Pulse Width	Minimum	20 ns
Tchpw	Minimum Positive CK Pulse Width	Minimum	6 ns
Tclpw	Minimum Negative CK Pulse Width	Minimum	6 ns
Tavch	Address valid to CK high	Minimum	4 ns
Tchax	CK high to address change	Minimum	1.0 ns
Tchdox	CK high to Data Output change	Minimum	2 ns
Tcd	CK high to data valid	Maximum	15 ns
Toe	Output Enable time	Minimum	0 ns
Toz	Output Disable time	Maximum	7 ns

20

"00000000000000000000000000000000"

**Table 8: Write Cycle for input RAM bank 2251**

NOTE: The RAM clocks for the input and output banks are limited to 10.24 MHz.

Symbol	Parameters	Condition	Nom
Tcc	Clock Cycle time	Minimum	20 ns
Tchpw	Minimum positive CK pulse width	Minimum	6 ns
Tclpw	Minimum negative CK pulse width	Minimum	6 ns
Tavch	Address valid to CK low	Minimum	4 ns
Tchax	CK low to address change	Minimum	1 ns
Twch	- WE low to CK low	Minimum	4 ns
Tchw	CK low to - WE high	Minimum	1 ns
Tdivch	Data Input valid to CK low	Minimum	4 ns
Tchdix	CK low to Data input change	Minimum	1 ns
Tchdov	CK low to Data Output valid	Maximum	15 ns
Tchdox	CK low to Data Output change	Minimum	2 ns

**Package for ASIC 2101**

**Package Dimensions and Pinout for one embodiment:**

The ASIC 2101 generates 5V TTL output levels and accepts 5V CMOS or 5V TTL input levels.

TTL Input Levels are defined as follows:

VIL max = 0.8 Volts

VIH min = 2.0 Volts

CMOS Input Levels are defined as follows:

$$\begin{aligned} V_{IL} \text{ max} &= 0.2 \cdot V_{DD} \quad (0.9V < V_{IL} < 1.1V \text{ over } 4.5V \text{ to } 5.5V \text{ VDD range}) \\ V_{IH} \text{ min} &= 0.7 \cdot V_{DD} \quad (3.15V < V_{IH} < 3.85V \text{ over } 4.5V \text{ to } 5.5V \text{ VDD range}) \end{aligned}$$

TTL Output Levels:

$$\begin{aligned} V_{OL} \text{ max} &= 0.4 \text{ Volts} \\ V_{OH} \text{ min} &= 2.4 \text{ Volts} \end{aligned}$$

### Functional Tests

Testing is broken down into a functional segment to verify device functionality and a scan segment to test for faults in the physical silicon. In one embodiment, the vectors are included in a test bench compatible with LSI Logic's tools. The functions to be tested in the device are listed below.

- a. **FFT/IFFT Operation** - Vectors are provided which are characteristic of the expected use in the system. The frequency domain vectors are passed through an inverse transform (with appropriate bit scaling) and the results stored. The vectors are then passed through a forward transform (with appropriate scaling), and this final result analyzed. These tests are performed for the 1024-, 512-, and 256-point transforms. There are 10 frames of data for each test. The test bench includes 1024 vectors for the 1024-point transform, 5120 for the 512-point transform, and 2560 for the 256-point transform. There are twice this number of vectors passed through the device to complete the test. The total number of test vectors for this test segment are about 36000.
- b. **FFT/IFFT Verification** - A single sinusoid is passed through the 128-, 64-, and 32-point transforms. Both forward and reverse directions are tested.
- c. **Bit Growth Tests**. Each bit-growth pin (2117-2118) is

exercised for the 1024- and 512-point transform in the forward and reverse direction.

- d. Power Down Tests - The device is placed in the middle of a transform, then powered down. The outputs are evaluated for correct state. The device is then asked to perform a forward and reverse transform to validate that the device can function after the reset.
- e. Overflow Tests - An overflow condition is induced, and the device evaluated for correct response (e.g. the overflow pin is actuated and the event does not cause an adder to wrap around). The test includes an overflow in the positive and negative direction.
- f. Reset - The device is placed into the middle of a transform operation, then reset. The outputs are evaluated for correct state. The device is then asked to perform a forward and reverse transform to validate that the device can function after the reset.

Figure 93 is a high-level block diagram of one embodiment of modem receiver 2402 as shown in Figure 71. The analog received signal-in is first processed by bandpass-and-down-convert block 2750. In one embodiment, the analog received signal-in is either 425 to 600 MHz or 550 to 770 MHz, and is converted by bandpass-and-down-convert block 2750 to a signal which is 100 kHz on both sides of a 18.432 MHz center frequency. In one embodiment, the tuning step size is 99 MHz. In one embodiment, analog-to-digital decimator system (ADDS) 2850 (in one embodiment, this is a Sigma-Delta decimator system that uses a Sigma-Delta analog-to-digital converter 2840) converts this band-limited signal into decimated I and Q quadrature signals, each 15 bits wide, which have a symbol rate of 288 K symbols per second (denoted 288 KBS), which are then processed by FFT block 2849. In one embodiment, FFT block 2849 is equivalent to FFT system

2100 of Figure 70. The outputs of FFT block 2849 are then processed by post-processing block 2990 into digital data out.

Figure 94 is a more detailed block diagram of modem receiver 2402.

Analog received signal in is fed to band-pass filter (BPF) 2740 which limits the input signal to either 425 to 600 MHz or 550 to 770 MHz. The signal is then amplified by variable-gain amplifier 2741, and mixed by mixer 2742 with a demodulator signal of either 627-802 MHz or 752-973 MHz generated by signal generator 2747 as controlled by phase-locked-loop block 2746. The mixer 2742 produces a difference-frequency signal centered at 202.752 MHz which is filtered by BPF 2743. The resultant signal is again amplified, this time by variable-gain amplifier 2744, and then mixed by mixer 2745 with a demodulator signal of either 221.184 MHz generated by signal generator 2749 as controlled by phase-locked-loop block 2748. The resultant signal is an analog signal centered at 18.432 MHz and having a 200 kHz bandwidth.

The resultant 18.432 MHz analog signal is then passed to an analog-to-digital decimator which, in this embodiment, is denoted as the components encircled by the dotted line referenced as Sigma-Delta decimator system (SDDS) 2850'. The exact mix of components which are included in SDDS 2850 can vary from embodiment to embodiment (i.e., where the dotted line for SDDS 2850' is drawn may vary).

In Figure 94, the 18.432 MHz analog signal is passed through BPF 2839 which is centered at 18.432 MHz. The 221.183 MHz signal is divided by divider network 2838 to produce a 73.728 MHz signal which drives Sigma-Delta converter 2840, and two 18.432 MHz signals (one of which is 90 degrees shifted from the other) which drive digital I/Q detector 2841. Sigma-Delta converter 2840 is any conventional Sigma-Delta converter, such as described in *An overview of Sigma-Delta Converters*, by Pervez M. Aziz et al., Vol 13, No. 1, IEEE Signal Processing Magazine, January 1996, which is hereby incorporated by reference. Sigma Delta converter 2840 achieves high resolution by oversampling the input signal at a frequency much above the Nyquist frequency, and by providing a negative feedback



path which uses an analog representation,  $y_a[n]$  of the quantized output signal  $y[n]$ . Figure 95 shows one such Sigma Delta converter 2840, having difference block 2860 that forms  $u[n]$  which is input  $x[n]$  minus feedback  $y_a[n]$ . Signal  $u[n]$  then is processed by discrete-time integrator 2863 and quantizer 2864, which is simply a comparator. By providing a sampling frequency which is high enough, the Sigma-Delta converter 2840 allows the use of a 1-bit quantizer to achieve high overall resolution.

Referring back to Figure 94, the 73.728 MHz quantized output of Sigma-Delta converter 2840 is coupled to digital I/Q block 2841. In one embodiment, digital I/O block 2841 is simply two 2-input AND-gates; one input of both AND gates is connected to the 73.728 MHz quantized output of Sigma-Delta converter 2840, the other input of the first and the second AND-gate is coupled to the  $0^\circ$  and the  $90^\circ$  18.432 MHz outputs of divider 2838. The outputs of digital I/Q detector 2841 are thus two serial streams, that represent I and Q quadrature signals respectively. The serial I and Q signals are then fed to coarse decimator and MUX 2842, which converts the two serial streams into a single N-bit-wide time-multiplexed I/Q parallel stream. In one embodiment, this time-multiplexed I/Q parallel stream is 10 bits wide. This time-multiplexed I/Q parallel stream is clocked at 2.304 MHz, and provides 10 bits of I followed by 10 bits of Q, thus having 1.152 million samples of I interleaved in time with 1.152 million samples of Q. This time-multiplexed I/Q parallel stream is then fed to I/Q demux 2843 which de-multiplexes the time-multiplexed I/Q parallel stream into separate I and Q streams clocked at 1.152 MHz each, and each 10 bits wide. These separate I and Q streams are then processed by three digital-processing blocks: DC-offset adjust block 2844 that digitally adjusts for DC (direct current) balance, DS\_gain adjust block 2845 that digitally adjusts the decimated signal gain, and DS\_mix block 2846 that digitally adjusts the phase.

The Sigma Delta decimator system (SDDS) 2850 is a N-bit A-to-D converter which generates a one-bit serial data stream having resolution and accuracy of N bits

(in one embodiment, 15-bit resolution is obtained; in another embodiment, the A/D has a 10-bit resolution with 9-bit linearity). SDDS 2850 is running on the clock generator 2749 divided to 73.728 MHz which oversamples the SDDS input signal in order that it only passes data at the  $18.432 \text{ MHz} \pm 100 \text{ kHz}$ , approximately. The following circuits 2841 - 2847 then take that 200 kHz of frequency that Sigma-Delta converter 2840 passes and shifts it down a base band, so basically it goes from 0 Hz to 200 kHz. This a relatively slow signal and then SDDS 2850 turns that into 10-bit parallel words. The Sigma-Delta converter 2840 is outputting a serial 1 bit stream, which is ANDed with two 18.432 MHz square waves to produce serial digital I and Q that are two 18.432 MHz gated square waves.

In one embodiment, the entire SDDS 2850 is integrated on a single VLSI chip using 0.8 micron BiCMOS 4S+ technology fabricated by IBM Corporation, with the analog circuits operating from a 5 volt supply voltage and the digital circuits operating from a 3.3 volt supply voltage. This single-chip implementation facilitates bit growth from stage-to-stage in order to prevent or reduce truncation or round-off errors. Thus 10-bit I and Q signals at the output of I/Q demux block 2843 are allowed, in one embodiment, to grow to 25-bits at the output of DS\_mix block 2846 through the digital processing of DC-offset adjust block 2844, DS\_gain adjust block 2845, and DS\_mix block 2846. For example, the N bits each of the I and Q data streams at the output of I/Q demux block 2843 grow a little to  $N^*$  bits after the digital processing of DC-offset adjust block 2844,  $N^{**}$  bits after DS\_gain adjust block 2845, and  $N^{***}$  bits after DS\_mix block 2846. Decimators 2847 and 2847' select one out of every four values from the I and Q data streams, respectively, thus producing a 288 kHz rate of 25 bits each for the I and Q streams. These then pass through scaling blocks 2848 and 2848' which scale each data stream to 15 bits, which are denoted I' and Q' and are coupled to 15-bit FFT 2849.

One consideration with the large number of signals on such a single-chip implementation is to minimize the number of different clock signals. In the embodiments described for Figure 94, for example, this is accomplished by running

a large number of blocks from a single clock, even though some of those blocks might be able to run off a different and slower clock.

#### Overview of Data-Delivery Architecture

Referring now to Figures 96 and 97, there is shown the overall architecture of a data delivery system 500 according to another embodiment of the invention. The data delivery system 500 provides high speed data access to the home 510 or office 512 (alternately referred to above in Figure 1 as remote units 46) over the HFC distribution network 11 using the infrastructure of the modem-based telephony transport system 10 described hereinabove. Figure 97 illustrates the integration of the data delivery elements in the HDT 12. The system allows users to have access to local content 520 and the Internet 530 through services available at the cable providers' premises or head end 32.

Among other things, system 500 provides: (1) user data access to the Internet 530 and local content on a head-end server 520 through the above-described access platform; (2) support for TCP/IP and transparent bridging at the data link layer using a Dynamically Adaptive Transport Switching (DATS) methodology (described below); (3) guaranteed, reliable symmetrical data service at transfer rates from 64 Kbps to 512 Kbps, in increments of 64 Kbps, for geographically dispersed individuals; (4) guaranteed, reliable, symmetric shared access to a 8.192 Mbps data pipe for geographically limited group of users connected in a routed configuration; (5) mixing of data and telephony within a single HDT 12; (6) network management for telephony, video, and data through integrated CMISE and SNMP; (7) routed service through a head-end server; and (8) use of HISU and MISU RF modem technology for transport.

Referring now to Figures 98-100, the Cable Data Modem (CDM) 535 for the system 500 preferably can take three forms: a stand-alone box called the Personal Cable Data Modem (PCDM) 540, a HISU add-in card called the Data Modem Service Module (DMSM) 550, or a MISU add-in card called the Data Modem Channel Unit (DMCU) 560. The stand-alone PCDM 540 has several

connector options; including standard F style coax cable, 10baseT Ethernet 542, and RS232 544. The DMSM add-in card preferably resides in an expansion slot on HISU 68 enclosure and will terminate a customer's computer with a 10BaseT Ethernet connection. The DMCU 560 add-in card will reside in a line-card slot of the MISU and will provide a 10BaseT Ethernet routed connection to four users. Customers without an integrated Ethernet port can add one to their system at minimal cost.

The system 500 provides connection to head end 32 services via the DATS methodology. The DATS methodology intelligently allocates bandwidth in the system to maximize both transport bandwidth and number of users while providing guaranteed bandwidth. Under TR303/V5 switching environments a central resource within the HDT 12 (described below) provides the intelligence to allocate bandwidth and efficiently manage transport load. This capability is implemented at the customer end by initiating a connection to the head end 32 when data is available to send. When the session is initiated, the head end 32 equipment determines the amount of bandwidth to be allocated to the subscriber as configured during pre-provisioning. The connection is maintained as needed and dropped when transmission is complete. During the time that the connection is maintained, bandwidth is guaranteed, providing the efficiency of packet switching and the guaranteed bandwidth of connection switching. All processing is performed through standard TR303/V5 call processing and therefore integrates completely with telephony provided over system 500. The architecture of the system 500 provides guaranteed bandwidth and latency to all supported users. As indicated in Figure 101, up to 24 concurrent users can be supported at 512 Kbps within a single 6 Mhz transport channel.

Using the DATS methodology, bandwidth is dynamically allocated to support a maximum of 240 users per 6 Mhz channel (e.g., see Figure 13) at 64 Kbps. Each user on the system is guaranteed the allocated bandwidth for the duration of the session. As such, the bandwidth and latency seen by the user

remains constant, independent of the traffic load, application type, or number of users. Assuming even distribution across all applications, the average frame size carried downstream is 378 bytes, while the average upstream frame size is 69 bytes.

The downstream latency from the HDT 12 is dominated by filling a 512 K data pipe, all other latencies in the system are insignificant. Therefore on average the latency to transfer a TCP/IP frame from the HDT to the customer is 5.9 msec. With Forward Error Correction (FEC) disabled, the average latency in the upstream direction, from the customer end to the HDT is also dominated by filling a 512 K data pipe. Therefore, on average the latency to transfer a TCP/IP frame from the HDT to the customer is 1.0 msec. Adding FEC will add up to 7 msec to the upstream latency.

#### Data Transport and Framing

The system 500 provides transparent bridging and data transport. A schematic representation of the data transport and framing is shown in Figure 102. Upstream data traffic begins at the customer end equipment 511 where frames are put out on the Ethernet link 542. The CDM 535 buffers the frames and encapsulates them with HDLC framing 570. The HDLC frame 570 is then sent out over the coax by a modem 101 as one or more of the possible 240 payloads on the coax. The frame, along with other data and telephony streams is converted at the ODN 18 and sent to the head end 32 over optical fiber. At the head end, the CXMU 56 demodulates and sends the stream to the CTSU 54, where it is routed to the appropriate LANU 580 as part of a MARIO data stream. The LANU 580 buffers the complete frame, stripping the HDLC framing prior to putting the frame on the local Ethernet. All routing for the Internet is performed by the head-end server 590. The routing services could also be incorporated in LANU 580.

In the case of downstream traffic, the head-end server 590 puts Ethernet packets on the LAN. One to potentially seven LANUs 580 examine the Ethernet address and selectively pick up packets destined for the customer-end

machines. Once the Ethernet packet has been accepted, the packet is framed in HDLC and sent, via a MARIO stream, to the CTSU 54 according to the routing established when the call was initially provisioned (explained further below). The CTSU 54 forwards the stream to the CXMU 56 where it is sent over the HFC distribution network 11 to the customer premises equipment 511. At the CDM 535, all packets received are passed on to the local Ethernet connection 542. In the case of a broadcast message, the LANU 580 at the head end will replicate the packet across all attached connections. This allows the bridge to support protocols such as ARP.

The system 500 will utilize the telephony error correction mechanism described above with respect to system 10. Under the telephony error correction scheme, forward error correction codes are generated for upstream traffic but not for downstream traffic. Forward error correction on upstream is generated at the ISU 100 (HISU 68 or MISU 66) and consumes the 10th bit of each DS0, thereby protecting each DS0 separately. The error correction can be disabled, but this is not recommended for data transport.

The error detection/correction processing occurs on the CXMU 56 and data is delivered corrected to LANUs 580, DS1Us 48, in MARIO streams. Therefore, the system 500 data architecture does not explicitly have to deal with error correction. The CRC of the HDLC frames provide for a level of error detection above the error detection/correction of the CXMU 56. Errors detected in the LANU 580 will be reported through the SNMP agent of the LANU 580.

#### Component Design

The data delivery hardware for the head end 32 equipment consists of the LANU 580 which interfaces with the HDT 12. The LANU 580 includes a board responsible for all data transport to and from the head-end server LAN 591. In that function, the LANU 580 operates as a point-to-multipoint connection that is responsible for concentrating up to 128 DS0s onto a single

Ethernet connection. The LANU 580 has the following features:

- (1) It sits in a DS1U 48 slot and is form/fit compatible with the HDT 12 backplane.
- (2) It provides 10BaseT Ethernet connection to head-end LAN.
- (3) It supports multiple LANUs 580 on a single LAN through a 10BaseT connection or fast Ethernet through hub.
- (4) It concentrates up to 128 DS0s into a single Ethernet Connection.
- (5) It supports time slot assignment and super-channel aggregation across four MARIO data streams independent of the CTSU 54.
- (6) It supports Dual HDT 12 LAN ports.
- (7) It maintains on-board FLASH for storage of operational code image.
- (8) It supports Nx64 service for super-channels up to 512 Kbps.
- (9) It guarantees time ordering of multi-channel calls independent of transport system.
- (10) It provides transparent bridging and broadcast of Ethernet frames between head-end LAN 591 and CDM 535.
- (11) It self-discovers Medium Access Control (MAC) addresses of CDM 535 and filters Ethernet frames with on-board CAM.

A high-level block diagram of the LANU 580 is shown in

Figure 103. The LANU 580 preferably employs a single processor, unified memory design. In this embodiment, the main processor on the LANU 580 is the Motorola MC68MH360 QUICC32 (581), running at 33 Mhz, which can deliver approximately 5 MIPS. The QUICC32 is actually a dual processor embedded controller that also contains a dedicated independent RISC processor communications called the CPM. The CPM along with Direct Memory Access (DMA) can access memory and move communication traffic without placing any performance burden on the main processor. In addition to acting as the host

processor, the QUICC32 can also act as a slave. The LANU 580 implementation is preferably based on 4 QUICC32s with one acting as master and the remaining three operating in slave mode. The combination of the four QUICC32s provide the majority of the system peripherals.

There are several types of memory 582 on the LANU 580. The first and largest is 8 MBytes of page-mode DRAM for the storage of LANU operational code. In addition to operational code, the DRAM also holds the routing tables, data buffers, and buffer descriptors needed to process data traffic. Second, the LANU contains 2 MBytes of FLASH memory to store an image of its operational code. On power-up, the image is loaded into DRAM. Finally, the LANU also contains 128 Kbytes of EPROM to hold boot code. The boot code will load the contents of the FLASH into DRAM on power-up. If greater performance is desired, fast SRAM can be added to hold the data buffers and buffer descriptors for the data traffic.

There are two sets of system peripherals on the LANU: those that are contained within the QUICC32s and those that are not. Regardless of location, peripheral control on the LANU 580 is managed by the System Integration Module (software) (SIM) of the master QUICC32. The SIM is responsible for managing generation of Chip Select (CS) and other control lines.

The most important of the system peripherals are the QMC HDLC controllers 586 located in the QUICC32s. The QMC of the QUICC32 can implement up to 32 HDLC controllers running at 64 Kbps and are used to process the data traffic in the MARIO streams from the CTSU 54. The QMCs 586 can also support super channels by aggregating multiple 64 Kbit channels into higher data rate channels. The QMC 586 is implemented in the CPM of each QUICC32 and with the addition of SDMA can autonomously move communications data to and from main memory 582. Operation of the QMC 586 is controlled by the master processor 581 through the use of buffer descriptors that reside in external memory 582.



In addition to the QMC 586, three additional HDLC controllers 587 are used to provide two 2 Mbps HDT backplane LAN ports and a 10 Mbit Ethernet port. As with the QMC 586, the master processor 581 controls the HDLCs through the use of buffer descriptors, but unlike the QMC, these data structures reside in dual ported RAM (DPR) that are internal to the QUICC32s.

There are two additional system peripherals that reside outside of the master and slave QUICC32s. The first is the Rate Adaptation/Ninth Bit Signaling/Time Slot Assigner logic block (RA/NBS/TSA) 588. The RA/NBS/TSA 588, as its name suggests has three functions. The rate adaptation function is responsible for converting the 2.56 Mbps, 9 bit data format of the input MARIOs to the 2.048 Mbps, 8 bit data format of the QMC 586 in both the receive and transmit directions. In both directions the RA is also responsible for managing any information that may be placed in the 9th bit.

The NBS is responsible for transmitting and receiving the data that is carried with each DS0 in the ninth bit. The ninth bit is used for two distinct functions. In the upstream direction the ninth bit carries information regarding the ordering of data within a multi-channel call. The signaling consists of a repeating number that indicates which time position the DS0 occupies in the multi-channel call. The format of the upstream ninth bit signaling is shown in Table 9.

Table 9: Upstream Ninth Bit Signaling

Bit	Contents	Description
1	"1"	Sync Pattern
2	"1"	Sync Pattern
3	"1"	Sync Pattern
4	"1"	Sync Pattern
5	"0"	Sync Pattern
6	"0"	Sync Pattern
7	D[8]	Order Number, Bit 8 [MSB]
8	D[8]*	Inv. Order Number, Bit 8 [MSB]
9	D[7]	Order Number, Bit 7
10	D[7]*	Inv. Order Number, Bit 7
11	D[6]	Order Number, Bit 6
12	D[6]*	Inv. Order Number, Bit 6
13	D[5]	Order Number, Bit 5
14	D[5]*	Inv. Order Number, Bit 5
15	D[4]	Order Number, Bit 4
16	D[4]*	Inv. Order Number, Bit 4
17	D[3]	Order Number, Bit 3
18	D[3]*	Inv. Order Number, Bit 3
19	D[2]	Order Number, Bit 2
20	D[2]*	Inv. Order Number, Bit 2
21	D[1]	Order Number, Bit 1
22	D[1]*	Inv. Order Number, Bit 1
23	"0"	Sync Pattern
24	"0"	Sync Pattern

The numbers can range from 1 to 128 with a 0 indicating that the DS0 has not been assigned a position. During a call, the processor 581 will monitor the ninth bit signals from all the channels and once the order has been established, the processor will configure the TSA function to order the channels. Once the order has been established, the processor will periodically monitor the ninth bit signaling to detect any changes in ordering (i.e., frequency hopping due to excessive errors).

In the downstream direction, the NBS is used to enable data transmission. Once the LANU 580 receives the ordering information for the

channels, the processor will enable transmission by sending a data pattern over the downstream ninth bit of the first DS0 of a multi-channel call. The format for the "Data Dial Tone" signaling is shown in Table 10.

**Table 10: "Data Dial Tone" Downstream Ninth Bit Signaling**

Bit	Contents	Description
1	"1"	Sync Pattern
2	"1"	Sync Pattern
3	"1"	Sync Pattern
4	"1"	Sync Pattern
5	"0"	Sync Pattern
6	"0"	Sync Pattern
7	CMD[15]	Command Bit 15, MSB
8	CMD[14]	Command Bit 14
9	CMD[13]	Command Bit 13
10	CMD[12]	Command Bit 12
11	CMD[11]	Command Bit 11
12	CMD[10]	Command Bit 10
13	CMD[9]	Command Bit 9
14	CMD[8]	Command Bit 8
15	CMD[7]	Command Bit 7
16	CMD[6]	Command Bit 6
17	CMD[5]	Command Bit 5
18	CMD[4]	Command Bit 4
19	CMD[3]	Command Bit 3
20	CMD[2]	Command Bit 2
21	CMD[1]	Command Bit 1
22	CMD[0]	Command Bit 0, LSB
23	"0"	Sync Pattern
24	"0"	Sync Pattern

Two commands are defined. The first is "Idle Sync" (bit pattern for CMD[15:0] is 0000 0000 0000 0000) is sent during idle times to synchronize the receivers. All idle syncs sent from the LANU 580 will be synchronized themselves to simplify the reception of order information in the upstream ninth bit signaling. The second command is "Data Dial Tone" (bit pattern for CMD[15:0] is 1000 0000 0000 0000)

and is sent once the order of the super-channel has been established. This command will instruct the CDMs to enable transmission.

The TSA is responsible for ordering and aggregating channels that are sent to the QMCs 586. Although TSA has the ability to aggregate up to 128 channels into a single data stream, most applications will aggregate multiple super-channels, up to 8 DS0s per super-channel, among four DS1-like channels feeding the four onboard HDLC controllers. For 1 to 32 DS0 data streams, the DS0s are aggregated and sent to the QMC 586 as a 2.048 Mbps serial stream. Within a single DS1-like stream, multiple super-channels can be supported by connecting to a single QUICC32 QMC 586. The TSA can also aggregate 128 DS0s into a single 8.192 Mbps data stream and connect it to a QUICC32 HDLC controller. In both cases the TSA is configured to insure that the time ordering of the data is maintained.

The final peripheral that does not reside in a QUICC32 is the Contents Addressable Memory (CAM) 589. The CAM performs memory accesses based upon data rather than address and is used to quickly determine whether an Ethernet frame should be accepted. The Ethernet controller 587a interfaces to the CAM 589 through glue logic and the reject input. When a frame is received that is not in the CAM, the CAM logic asserts the reject control line and the received portion of the frame is discarded. The buffer depth of the Ethernet controller is set so that no memory accesses are generated on rejected frames. The CAM is available off the shelf from MUSIC semiconductor.

The LANU 580 sits in a DS1U 48 slot at the HDT 12 and is form and fit compatible with the DS1U 48 to be backplane compatible. The LANU 580 has four major system connections: MARIO 592, Ethernet 593, HDT LAN 591 and clocking (input not shown in Figure 103). The four MARIO connections connect the LANU to the CTSU over the backplane. The four MARIO connections provide up to 128 bi-directional, 64 Kbit channels. The Ethernet connection connects the LANU to a local 10 Mbit, 10BaseT LAN. The 10BaseT connection will take the place of a single T1 connection of a DS1U. The connection is

terminated by wire-wrapping the 10BaseT to the backplane and routing to a patch panel. For call provisioning and other system functions, the LANU 580 connects to two common 2 Mbps HDLC LANs on the backplane. Finally, system clocking is provided by an external source such as a Building Integrated Timing Supply (BITS).

#### Functional Description of Upstream Communications

For upstream traffic, the HDLC framed data available on the MARIO interface passes through the rate adaptation and TSA block 488. In this block, the 2.56 Mbps MARIO interface is rate adapted down to 2.048 Mbps. As part of rate adaption, the ninth bit of each DS0+ of the MARIO stream is stripped and sent to the NBS logic. The ninth bit carried an order number that is used to time order the DS0s in multi-channel calls. Once the order numbers are established, the processor 581 configures the TSA to re-order the multi-channel calls and target the super-channel to a QMC 586. For super-channels composed of 32 or less DS0s, the call is placed in a single 2.048 Mbps data stream, along with other calls and sent to a QMC 586. For 128 DS0 calls, the DS0s are placed in a single 8.192 Mbps stream that is target to a QUICC32 HDLC controller 586 configured for 8.192 Mbps HDLC. Whether targeted to a 2.048 Mbps QMC or 8.192 Mbps HDLC, the frames are accumulated and transmitted on the local Ethernet LAN.

#### Functional Description of Downstream Communications

For downstream traffic, data on the LAN is filtered according to the destination MAC address. If the MAC address is in the Content Addressable Memory (CAM) 589, the LANU 580 will accept the Ethernet frame. Once the frame is accepted, the LANU 588 accesses a routing table in memory 582 to select the appropriate MARIO slot for transport. The frame is then scheduled for transmission and the HDLC controller 586 takes over. In the downstream direction for 32 DS0 or less calls, the HDLC controller 586 is responsible for creating the MARIO stream and encoding the data into HDLC format. For 128 DS0 calls, the 8.192 Mbps HDLC data stream is split among the four MARIO interfaces (A-D).

The time ordering and aggregation in the downstream direction is controlled by the TSA. After the data passes through the TSA, the ninth bit signaling information is added to indicate that data transmission is enabled. At the same time that the ninth bit signal is added, the data stream is rate adapted up to the 2.56 Mbps rate of the MARIOs.

For data movement the on-board memory 582 provides the switching and buffering necessary to move traffic from the Ethernet to the MARIO and vice-versa. The data switching provides a virtual circuit based upon routing information constructed during call provisioning that maps the MAC address to MARIO time slot. The LANU 580 constructs this routing table from information provided during call provisioning and by examining the source MAC address of Ethernet frames received from the customer end. The MAC addresses supported by the LANU are then placed in the CAM 589 where they can be used to filter Ethernet traffic.

In addition to the MARIO and Ethernet interfaces, the LANU 580 also supports two HDT backplane LAN 591 interfaces at 2.048 Mbps. The HDT LAN is maintained by the SCNU 58 at the head end 32 and is used to pass configuration information to the LANU 580. The combination of LANU 580, call processing, and pre-provisioning will provide the MARIO route (time slot in the MARIO data stream) to all customer devices.

#### Customer Premises Equipment (CPE)

The CPE of the system 500 comes in different implementations. Although the basic data framing and transport remain the same across all the implementations the underlying modem technology and form factor differs. There are currently three general types of CPE defined: HSU based, MISU based, and Stand-alone (a variant of the HISU based implementation). Each of these implementations are discussed separately below.

### Data Modem Service Module

A simplified block diagram of the DMSM is shown in Figure 104. DMSM 550 supports up to 8 DS0's for data traffic. Data is interfaced to the HISU 68 through a Service Growth Module (SGM) interface that implements the bridge router functions of the data connection. The bridge/router is based upon a 68302 embedded processor 605 with 2 MBytes of DRAM 604 and 256Kbytes of FLASH memory 606. One embodiment of the customer interface is a 10BaseT, 10 Mbit Ethernet connection 602. The interface between the bridge/router and the HISU modem is a Super-LUIGI (SLUIGI) interface operating at 2.048 Mbps. Data in the upstream direction is accepted by the Ethernet connection and relayed to the HISU 101 modem over the SLUIGI interface as HDLC encoded data. In the downstream direction, data from the SLUIGI interface is passed to the HDLC controller 607 and then sent out the Ethernet connect through Ethernet controller 609. System memory consists of 2 Mbytes of DRAM 604 which contains the operational code that implements the bridge/router functions as well as an SNMP agent. The 256Kbytes of FLASH 606 is used for storage of the operational image and can be updated with TFTP transfers.

In the upstream direction, the HISU interface logic 608 is responsible for generating the ninth bit signaling information for each DS0 and interfacing with the SLUIGI stream of the HISU. The 2.048 Mbps serial data from the HDLC controller is stuffed in the appropriate time slots of the SLUIGI interface. In the case of 64Kbit traffic, all HDLC data is placed in a single SLUIGI time slot. In the case of multi-channels calls, all HDLC is placed in adjacent SLUIGI time slot receiving the first byte of HDLC data. In addition, the interface logic will generate an order number for each DS0 with "1" assigned to the first DS0 and "8" assigned to the last DS0, "0" is reserved to indicate that the position of the DS0 has not been set. In the event of frequency hopping, the interface logic will continue to number the DS0 time-slots in the order that they are received. For example, under normal circumstances DS0 ninth bit signaling would tag the time slots as "12345678-" for

an eight channel call. If DS0 4 is lost due to ingress, then the new order would be "123-45678." This signaling information is used at the head-end to reorder the DS0s independent of the frequency hopping and transport services.

In the downstream direction the HISU interface logic 608 is responsible for taking data from the SLUIGI interface from the RF modem 101 and giving it to the HDLC controller. In addition, the HISU interface logic monitors the ninth bit signaling information of the first DS0 to detect the "Data Dial Tone" sequence. The "Data Dial Tone" sequence is sent by the head-end to enable data transmission. During the call provisioning process the HISU interface logic sends the ninth bit ordering information as soon as the SLUIGI time slots indicate that they can accept data. It is not until the HISU interface logic gets a positive acknowledgment through the "Data Dial Tone" that data is sent to the head-end in the upstream direction.

#### Data Modem Channel Unit (DMCU)

A simplified block diagram of a DMCU 610 is shown in Figure 105. The DMCU supports up to 128 DS0s for data traffic. The data interface to the MISU is a specialized channel that sits on the MISU backplane.

The basic design of the DMCU 560 is very similar to the design of DMSM 550. The most notable difference is the interface to the RF portion is two 8.192 Mhz serial channels. This allows the MISU interface to support a symmetrical 8.192 Mbps Ethernet connection. Because of the higher throughput the MISU is based on the MC68360 614 that can support both the 8.192 Mbps HDLC connection as well as the 10 Mbit Ethernet. In front of the 10 Mbit Ethernet interface 612 is a router 616 that allows four users access to the 8.192 Mbit data connection. The router 616 design insures security for all connected users. The design contains 2 Mbytes of DRAM 618 and 256bytes of FLASH 619. Like the HISU design, the FLASH can be remotely updated with TFTP.

The DMCU 610 has an equivalent interface function that moves data from the HDLC controller 611 to the RF modem and works in a very similar

T06020"42ET0660



way to the DMSM 600 with the exception that the MISU modem interface is formatted as two SLUIGI streams that are clocked at 8.192 Mbps. Between the two SLUIGI streams; 128 DS0s can be carried between the HDLC controller 611 and the modem. The MISU interface logic 613 is responsible for buffering and then sending data over the dual S-LUIGI interface to and from the RF modem. In the upstream direction the MISU interface logic 613 generates order numbers for each of the 128 DS0s over the ninth bit of the DS0+. The order numbers generated on the MISU work in the same way as they do on the HISU. In the downstream direction, the MISU interface logic 613 is responsible for moving data from the S-LUIGI streams to the HDLC controller 611. The MISU interface logic also monitors the ninth bit data stream from the first DS0 to detect the "Data Dial Tone" that enables data transmission. Ethernet controller 617 moves the data to the router 616.

#### Stand-alone Data Modems

The stand-alone data CDMs are based upon the HISU design. In the stand-alone designs, the RF modem 101 of the HISU is tightly integrated with the bridge/router design. Like the DMSM 600, the stand-alone supports from 64 K to 512 K. The interfaces are identical with several options: standard "F" style connector to the cable, 10 Mbit Ethernet and RS232 connection to the customers equipment.

#### System Software and Call Processing

In the TR-008/V2 system, calls are provisioned and nailed up at time of installation. Under this scenario an operator at the head-end 32 is responsible for determining the MARIO configuration and transfer rates (64 K to 512 K). The DATS methodology of present invention utilizes TR-303/V5 call processing to provide dynamic allocation of bandwidth. To maintain the telephony oriented architecture of the access platform of the present invention, the LANU 580 takes on responsibility of a limited subset of the Central Office (CO) functions.

5

10

15

25

of bandwidth per data connection. For example, a subscriber may sign up for 512K of bandwidth. Upon call setup or connection, the DAT methodology thus assigns each user the number of channels required to achieve the subscriber's bandwidth. However, in certain cases, the transport system 500 may not have the necessary bandwidth to allocate to a subscriber their normal subscribed bandwidth. Under these circumstances, the subscriber is allocated a lesser amount of bandwidth, for example 64K of bandwidth. By dynamically adjusting the amount of bandwidth assigned at call setup for each call, the mLANU can maintain a minimum level of bandwidth for each subscriber. However, existing constraints prevent small decremental bandwidth re-allocations. Constraints existing within the system 500 as described above will enforce the halving of bandwidth of some subscribers to accommodate additional subscribers, rather than by means of a more evenly distributed loading. As shown in Figure 107, as the number of subscribers increase (on the horizontal axis) any given user's allocated bandwidth will halve at a certain user density. The spread around the average indicates that some users will necessarily lose half their existing bandwidth earlier under loading than will others. The constraining factor in preventing a more equitable burden is the window nature of the present HISU 68 RF tuning. As noted above, the HISU tunes to one of 24 IOC channels spread throughout the 6MHz cable channel and has access to five payload channels above and five payload channels below the selected IOC frequency. It cannot borrow payload channels outside this window of ten channels, so therefore there is no way for the "25th" user to borrow just one channel each from seven other users. It can only sit on top of one of the existing IOC payload windows and takeover half of the window bandwidth. At that point, 23 users would be granted 512kbs bandwidth and two users would each get 256kbs bandwidth. The "26th" user would result in 22 users with 512kbs and four users with 256kbs, and so on. This general pattern is repeated at a load of 72 users and 120 users. (There is a discontinuity in the pattern from 48 to 72 users due to the previously unused two of ten DSO's per window being pressed into two DSO 128 kbps service.) The graph of

**S**

10

15

20

25

communications are enabled and data transmission begins.

A session is terminated at the customer end when no data is available for transmission by generating an "On Hook" message. The call processing sequence for an "On Hook" message is shown in Figure 110. When the CDM terminates the connection, an "On Hook" message is sent over the IOC to the CXMU 56. The CXMU 56 in response sends an "On Hook" message, identifying the CRV, to the CTSU 54. The CTSU 54 then sends a "Tear Down" message to the mLANU 584 over the TMC. At the mLANU 580, the connection is deleted from the connection database and then released. If the connection is not on the mLANU, the mLANU will send a "Release Channel" message to the target LANU 580 and also will send a "Release Cross Connect" message to the CTSU 54. The CTSU 54 will release the cross connects used for the connection and then send a "Release Bandwidth" message to the CXMU 56. At the CXMU 56 the mapping between tones and DS0s is lost and the connection is lost. When the connection is lost, the CDM will lose the "Data Dial Tone" in the ninth bit signaling of the first DS0 of the call.

The LANU 580 can also be configured to bring up connections to customer-end equipment. This allows for notification of incoming e-Mail and personal Web pages without tying up idle bandwidth. To do this, the master LANU in each system will maintain a mapping between the MAC address for each data element in the system and cross that with the CRV. Then an Ethernet packet is put on the head-end LAN 591, and a LANU 580 will read its MAC address and determine whether the connection is up to the device. If the connection is up, the packet will be forwarded over the HFC transport. If the connection is not in place, the receiving LANU 580 will generate a connection request to the mLANU. The mLANU will then signal the transport system over the TMC to bring up a connection to that device using the IOC. The receiving LANU 580 will then send the data once the connection has been established.

106020-12370660

The LANU software 620 is responsible for the all major function of the data concentration of the head-end equipment. A simplified schematic diagram of the LANU software is shown in Figure 111. The software 620 of the LANU consists of three major components: bridging 621, HDLC LAN manager 622, and data IDT 623. All three tasks will operate as applications on top of the embedded controller operating system "pSOS" kernel in the processor 581. The pSOS kernel will provide the base for the multi-tasking operation of the software 620.

The most important task to the actual transport of data will be the bridging task. The bridging task has several functions. First, the task will be responsible for providing the virtual switch between the MARIO and Ethernet interfaces. The task will be implemented as in "interrupt on receive" task that will execute at interrupt level. At either interface, an interrupt is issued when an entire frame has been received and stored in buffer memory (Figure 103, 582). During the interrupt service routine, the packet will be handed off by modifying the associated buffer descriptor after looking up the routing in the bridging table (stored in memory 582). For upstream traffic (HDLC to Ethernet), the source of the first packet will be read to discover its MAC address. This address will be added to the bridging table and written to the Ethernet CAM 589 for filtering.

A second function of the bridging task is the creation and maintenance of the bridging table. The bridging table will match the MAC address of the CDM 535 with the MARIO DS0s so that data can be moved between the Ethernet and HDLC. During call processing, the DS0s that are allocated to the call will be identified by the mLANU 580 through backplane LAN (591) Messaging and installed in the bridging table. As described above, when the first frames begin to flow from the CDM, the source MAC address will be identified and the table entry for the CDM will be complete. At this point data will flow in both directions. Once

the MAC address has been discovered, the bridging table entry will remain intact until the connection is terminated by the CDM.

A third function of the bridging task is maintenance of an SNMP agent. SNMP traffic will be handled and processed from the Ethernet port. The agent will support a standard MB-II information database for transparent bridging. In addition, objects will be added to the MIB that are specific to the data architecture to facilitate CMISE-SNMP integration and different billing options.

Finally, the bridging function may support link-layer encryption/decryption on the bridged data. Encryption/Decryption may be software only or hardware assisted depending upon the desired performance of the system. In either case, this function will execute as an application on top of pSOS.

Another component of the LANU software 620 is the HDT LAN manager 622. The HDT 12 LAN 591 is used to communicate system messages between the elements of the HDT 12. During pre-provisioning, the SCNU 58 will communicate system parameters such as CRV, IDT ID, and number of channels accessible by the CDM to the mLANU 580. These parameters will be used in the construction of the call provisioning table resident on the mLANU. During call provisioning, the mLANU will examine the provisioning table for available DS0s and use the HDT LAN 591 to set up MARIO configurations on other LANUs 580.

Another important function of the HDT LAN manager software 622 is support for field software upgrades. During download, the LANU 580 will take the image from the HDT LAN and store it in on-board FLASH memory. Aside from the SCNU 58, the LANU will be the only board in the HDT 12 that will load its image from its own FLASH on power up. As such, support for image download from the SCNU 58 Ethernet port will need to be added to the SCNU software. A final function of the LAN manager software 622 is to provide the network management access to the SNMP environment of the LANU.

The final major task of the LANU software 620 is the data IDT

106020-120660

5

10

15

25



interconnect so that the order within a multichannel call is maintained. In addition, the data IDT software 623 will communicate the state of the connection to the bridging task to open up the data pipe.

Another important function of the data IDT software 623 is to provide the information needed to provide billing and other accounting functions. As an IDT like function the data IDT will also support standard CMISE objects.

#### CDM Software

Software provided in a CDM 535 (executing on the local processor 605 or 614 and represented by such elements) will provide the same types of functions. The major function of the CDM software is to provide the bridge/router (brouter) functionality at the customer-end. In supporting the brouter function the CDM software supports IP routing, PPP, and SLIP. As part of IP support the CDM supports TFTP for downloading new code images. The CDM also supports a standard SNMP agent with a full MIB-II information base. Preferably, the software executes on either the 68302 (605) or 68360 (614) processor.

The control of the modem portion is with the standard MISU or HISU software running on a Motorola 68HC11. This code supports all the alarm conditions and communications set out above for IOC communications. The interface between the RF modem and the brouter is preferably a hardware only implementation.

#### Network Management

Network management of the data architecture of system 500 is preferably provided by both CMISE and SNMP. The CMISE portion of network management will be responsible for the transport mechanism for data, while SNMP will be used for data network oriented management. In this environment, SNMP is an overlay to the CMISE environment.

As with all telephony services, the data architecture will depend

**S**

10

## 15

20

25

channels of the telephony transport system 10. In addition to the downstream modulator, the ASMU 622a concentrates the return channels by interfacing with the HDT 12. Upstream traffic is carried over as a single DS0 and integrated with the downstream transport on the ASMU 622a. The connection from the ASMU to the head-end services is 10BaseT, but higher capacity industry standard connections are also possible.

The ASMU 622a sits at the head end 32, but not in the HDT 12. The function of the ASMU 622a is to integrate the upstream path for up to 400 DS0s (Configurable from 64 Kbps to 512 Kbps) and a 30 Mbps shared downstream. Each LANU 580 will generate an 8.2 Mbps HDLC stream that contains the 64 Kbps Ethernet packets from all the users that are attached to LANU through the transport system. On the ASMU 622a, up to four of these, are aggregated, and sent out to the head services over 10BaseT Ethernet. In the downstream direction, the data on the 100BaseT is filtered on the ASMU, and those packets destined for the customer end products are accepted and then modulated onto the 30 Mbps shared medium.

In order to register a modem, the customer-end modem sends out an IP packet to identify itself. This causes the LANU 580 to assign an HDLC address that is mapped to the MAC address of device. This information is passed to the ASMU 622a so that the HDLC address can be used by the modulator over the HDT 12 backplane LAN 591. The HDLC address and frequency for the tuner is also sent to the customer-end over the downstream telephony path and registered at the customer end. This address is then used by the customer-end equipment to filter the 30 Mbps downstream channel.

One advantage of the asymmetrical system is that a relatively large number of casual users (300+) can be supported by a single multi-megabit downstream transport, with an optimal amount of upstream capacity. The implementation of the downstream matches the downstream of other cable data modems in use and additionally provides superior, high capacity upstream. Since casual users place lesser demands on the network (peak utilization is lower than that

of business), users can be concentrated on the return channel, thus lowering head-end 32 costs.

The upstream channel in asymmetrical applications is still important due to the nature of the acknowledge protocol of TCP/IP, where blocks of data sent in the downstream must wait on an acknowledge message from the receiver before subsequent data blocks are sent. If the upstream channel is too small, the downstream channel will stall, reducing the utilization of the downstream bandwidth. By guaranteeing a minimum of 64Kbps to each user, the asymmetrical system can deliver greater than 1 Mbps sustained to each user, matching the capacity of most residential computer equipment. Another advantage is the superior security of OFDM in the upstream. Unlike other shared upstream modem products currently available, the asymmetrical system herein described prevents information, such as bank accounts and credit card numbers exchanged during on-line Internet shopping, from being "seen" by other modems on the network.

#### Summary of Data Delivery System Advantages

Thus, the symmetrical embodiment of system 500 provides many options for the delivery of data services over HFC distribution network 11 to the residence or business. The DMSM 550 provides from 64 Kbps to 512 Kbps access to head-end resources over a 10BaseT connection or RS232 (64 Kbps service). The service is symmetrical (same data rate upstream and downstream), non-shared and dedicated to each user, providing a guaranteed level of service. As an add-in card to the HISU, the DMSM 550 provides complete transport integration with telephony, supplying high-speed data and two POTS lines to the residence.

The PCDM 540 provides the same data transport capabilities as the DMSM 550 in a standalone configuration, packaged in a traditional modem housing. This implementation is ideal for premises or installations where telephony is not deployed.

The DMCU 560 is an MISU channel unit that provides higher

data rates than either the DMSM or PCDM-512K. The DMCU 560 router manages four subscribers who share up to 8.192 Mbps of symmetrical bandwidth. The router implementation guarantees that all four subscribers on the DMCU 560 have private connections. The DMCU 560 works well for multiple dwelling installations for Internet access and small business connections where symmetrical, non-shared data access is required.

At the head end, the LANU 580 provides the concentration of up to 100 DS0s in flexible combinations of various data rates, from 641 Kbps to 512 Kbps for residential, and up to 8.192 Mbps for business applications on a single, industry standard 10BaseT connection. An HDT 12 can be configured with up to seven LANUs, concentrating up to 700 DS0s. In addition to the industry standard transparent bridging function, the LANU also provides the intelligence for the dynamic, adaptive allocation of bandwidth capacity to optimize transport during times of heavy loading. This capability enables an HFC service provider using system 10 to mix residential and business data services in a single 6 Mhz channel without compromising the quality of service for business connections during peak Internet access times. Dynamic-allocation allows the customer units to efficiently utilize the data transport by dropping connections at times of no traffic and re-establishing them when data is ready to send. Each time a connection is established the LANU 580 will allocate bandwidth of up to a maximum of 512 Kbps, depending upon the network load, with a minimum of 64 Kbps. Finally, the LANU collects detailed traffic statistics that can be used for a variety of billing methods, for instance bill by connect time.

System 500 is particularly effective in meeting the special needs and higher expectations of business applications. Businesses tend to require a higher level of upstream signaling in order to support applications such as telecommuting and videoconferencing. Most cable data modem network architectures can provide only limited upstream capacity, but ADC is able to, offer a very high capacity upstream due to the efficiency of OFDM and frequency agility.

5

10

15

20

25

**System 500 provides a completely scalable data architecture by dynamically allocating bandwidth for data traffic through its utilization of a subset**

of standard TRJ03/V5 call processing software. This system capability gives HFC service providers the flexibility to tailor the configuration of head-end resources to satisfy the diverse needs of their subscriber base. Subscriber services can be provisioned at the head end as symmetrical fixed, symmetrical variable, or asymmetrical services. As the subscriber mix changes or subscribers upgrade service, head-end resources can be re-provisioned to meet the new requirements. For example, users can be easily reconfigured to upgrade from 64 Kbps service to 512 Kbps or even from asymmetrical to symmetrical. For capacity planning, data bandwidth is allocated as a number of DS0s to potential users with a single HDT supporting up to 720 DS0s. The number of users supported is then a function of service level (number of DS0s) and concentration ratio (number of users per DS0).

To ensure that service providers have an effective tool to manage their cable data networks, system 500 offers an integrated data/telephony network management solution. Data management is based on industry standard SNMP agents and MIBs (management information bases), which are then combined into an integrated data/telephony network management environment. Integration of data delivery and telephony into a single network management system has several advantages:

- (1) Symmetrical data, asymmetrical data, and telephony elements can be managed by the same element manager.
- (2) Less support staff is required.
- (3) Better integration with billing.
- (4) Better fault isolation.
- (5) Lower Mean Time To Repair (MTTR).

Thus, system 500 provides a single, integrated system that can meet the diverse needs of potential subscribers, from casual Internet browsers to high-capacity business users. The integrated solution gives HFC service providers a single point of network management that results in reduced support costs, reduced staffing costs, and shortened time to turn-up new services. Finally, the OFDM

technology of system 500 provides data, video and telephony services in a bandwidth-efficient system that reduces the demands on a very valuable commodity for HFC service providers spectrum.

### ATM Embodiment

The system 500 of the present invention can also be configured to carry data from an Asynchronous Transport Mode (ATM) network. As shown in Figures 114 and 115, system 10 or 500 of the present invention is modified to include an ATM multiplexer/modulator 650 which can receive ATM data from an ATM network 652 and modulate it onto the HFC network. In one preferred embodiment, digital video data is delivered over ATM network 652, multiplexed and modulated using multiplexer/modulator 650 onto the HFC network in RF digital OFDM format on assigned data and/or telephony channels between the head end and a subscriber, as for example described above with respect to system 10 or 500. A digital set top box 654 receives the digital video, formatted for example in 4.0 Mbps MPEG or equivalent, and converts it to video for display on a television 656. A return path to the HDT 12 over a telephony or data channel allows for interactive digital video. A video server 658 and ATM switch 660, feeding the ATM multiplexer/modulator 650, is shown in Figure 115.

### Embodiment of control aspects of telecommunications system

#### Channel Manager

In one embodiment, communication system 10 of Figure 1 includes channel manager 900 of Figure 59 to control various aspects of the dynamic allocation of channels to ISUs 100. For example, channel manager 900 assigns each ISU 100 to a subband, allocates channels in the subband to an ISU to complete a communication link, and monitors the channel to detect and avoid use of corrupted channels. Channel manager 900 implements further functions as described below to coordinate the use of the channels in a 6 MHz transmission channel to ISUs 100.



Channel manager 900 may comprise software executed by a processor resident in each CXMU 56 of each HDT 12. Channel manager 900 receives events from board support software 902, IOC and modem communicators 904, ISU ranger 906, and administrator 908. Channel manager 900 also sends messages to IOC and modem communicators 904 for allocation or reallocation of channels. Channel manager 900 uses two types of channels to communicate control data to the ISUs. First, channel manager 900 broadcasts control data over the IOC channels to the ISUs. The control data on the IOC channels contains an identification signal that indicates the ISU to receive the control data. Further, channel manager 900 uses an ISU demand link channel, referred to as an IDL channel, for non-time-critical transport of data between head end 32 and an ISU when the data is of a size that would benefit from a transmission channel with more bandwidth than the IOC. Typically, the data rate for the IOC channel is 16 Kbps and the data rate for the IDL channel is 64 Kbps due to the amount of data contained in each package or frame. Typically, control signals contain four data bytes or less per frame or package. The IDL channel is used to transmit data packages that are larger than this. For example, the IDL channel is used to download software to an ISU, provision a channel unit, transmit future channel unit functions, or transmit protocols. In one embodiment, HDT 12 only implements one IDL at a time. The IDL channel is described in more detail below.

#### Subband Assignment and Channel Allocation

Channel manager 900 is responsible for assigning an ISU to a subband and for allocating payload channels for communications links to the ISU. Appropriate selection of subband and payload channel improve the performance of communication system 10. Channel manager 900 further monitors the channels and reassigns subbands and reallocates channels as necessary to maintain acceptable communications links between head end 32 and ISUs 100.

#### Subband assignment

Channel manager 900 selects a subband for an ISU in several circumstances:

during acquisition, when an HISU is assigned to a subband that has insufficient payload channels to meet a request, and during an HISU IOC timeout event. An IOC timeout event occurs when acknowledgments are not received by channel manager 900 from an ISU within a specified time period. With a timeout, it is assumed that the downstream communications to the ISU are still in tact even though the upstream communications have become corrupted due to noise or collisions. Thus, a message on the IOC to retune to a new subband is assumed to reach the ISU despite the lack of an acknowledgment.

In each case in which an ISU is assigned to a subband, channel manager 900 uses various criteria to select the subband for an ISU. Figure 62 is a flow chart that illustrates one embodiment of a method for assigning an ISU to a subband. According to this method, channel manager 900 first selects a subband. Channel manager 900 then determines whether addition of the ISU to the subband would provide an acceptable load on the IOC channel. For example, channel manager considers the number of ISUs assigned to a subband. Further, channel manager considers the type of ISU and the likely load that the ISU will place on the IOC channel. By considering these factors, channel manager 900 can selectively distribute the load on the IOC channels so as to facilitate timely communication of control data to and from the ISU. This also allows channel manager 900 to evenly distribute the ISUs over the available subbands such that a like number of ISUs occupy each subband. Channel manager 900 also weighs the number of available channels within the subband and their transmission quality as recorded in the tables of channel manager 900. Channels with longer low-error rate histories will be used first. Channels previously marked bad and reallocated for monitoring will be used last. Based on these criteria, channel manager selects a subband for each ISU.

Figures 63, 64 and 65 are frequency spectrum diagrams that illustrate initial assignment of HISUs and MISUs to various subbands in a 6 MHz transmission channel. These Figures show that channel manager 900 attempts to evenly distribute the ISUs across the transmission channel. As depicted in Figure 63,

channel manager 900 begins assigning subbands at the middle of the 6 MHz transmission channel. Channel manager 900 then moves out toward the ends of the transmission channel. For example, the first HISU is assigned to subband number 12 and the twenty-fourth HISU is assigned to subband 0. It is noted that more than one ISU can be assigned to a subband. As depicted in Figure 64, channel manager 900 initially assigns the first MISU to subbands 0 through 12 and the next MISU to subbands 11 through 23. As depicted in Figure 65, when HISUs and MISUs are assigned to the same subbands, channel manager assigns the subbands so as to evenly distribute the ISUs over the available subbands. It is noted that the factors listed for use in selecting a subband are shown by way of example and not by way of limitation. Other factors can be added and the weight given to each factor can be adjusted without departing from the spirit and scope of the present invention.

#### Channel allocation

Figure 60 is a flow chart that illustrates one embodiment for a method for allocating payload channels in a subband by channel manager 900. Channel manager 900 attempts to maintain an acceptable distribution of bandwidth within a subband to reduce the need for reallocation of payload channels within the subband. Further, the goal is to allocate channels appropriately across the 6 MHz transmission channel to avoid having to reallocate channels that are currently in use. A channel can be allocated to an ISU only from the available channels in the subband to which the ISU is assigned.

Channel manager 900 receives a request for allocation of a payload channel from either the SCNU 58 or CTSU 54. At block 912, channel manager 900 decides whether sufficient payload channels are available in the current subband to fulfill the request. If sufficient channels are available, the method proceeds to block 914 and determines whether one of the available channels is the IDL channel. If the IDL channel is not one of the available channels, channel manager 900 allocates a channel for each channel requested by CTSU 54 or SCNU 58 at blocks 916 and 918. Channel manager 900 selects the channels based on several criteria that increase the

likelihood of achieving a connection with acceptable quality levels. For example, channel manager 900 can use the method shown in Figure 61. According to this method, channel manager 900 begins the selection process by identifying available payload channels that are located toward the center of the 6 MHz transmission

5      channel. Typically, channels that are nearer to the edge of the 6 MHz channel exhibit higher bit error rates than the channels that are closer to the center. Further, channel manager 900 can also consider limitations of the ISU and the requested service in selecting a payload channel. For example, the ISU may be preset for use only with odd or even payload channels. This information may be included in a

10      ROM on the ISU and provided to the channel manager when channel allocation is requested or during acquisition. Further, channel manager 900 uses data on the quality of transmissions over the identified channels stored in tables in channel manager 900 to determine which available payload channels have an acceptable error history, e.g., bit error rate. Other appropriate criteria can be used in channel

15      selection that also tend to increase the chances of producing a connection with acceptable quality. Based on these criteria, channel manager selects a payload channel to allocate to the ISU.

If, at block 914, channel manager 900 determines that one of the available channels is the IDL channel, channel manager 900 deallocates the payload channel allocated to be the IDL channel at blocks 920 and 922 due to the lower priority of

20      communications over the IDL channel.

If, at block 912, channel manager 900 determines that sufficient payload channels are not available in the current subband, channel manager 900 determines whether the request is for an HISU 68 or an MISU 66 at block 924. If the request is

25      for an MISU 66, channel manager 900 sends a message to the requestor that the request has failed at block 926.

If, at block 924, channel manager determines that the request is for an HISU, then channel manager 900 selects a different subband at block 928 by weighing the criteria as described above with respect to selecting a subband. Channel manager

00004370001  
00004370001

5

10

## 15

20

## 25

Channel monitor 900 monitors the payload channels for errors to help in determining which channels are acceptable for transmission for specific services. One input to channel manager 900 is parity errors which are available from hardware per the DS0+ channels; the DS0+ channels being 10-bit channels with one of the bits having a parity or data integrity bit inserted in the channel as previously discussed. The parity error information on a particular channel is used as raw data which is sampled and integrated over time to arrive at a quality status for that

channel. In one embodiment, parity errors that are detected on downstream payload channels are communicated to head end 32 over an associated upstream channel. When the error is detected in the downstream transmission, the parity bit for the upstream transmission is corrupted by intentionally setting the parity bit to the wrong value to indicate the incorrect parity in the downstream transmission path. Thus, the ISU informs the head end of errors in the downstream path.

To monitor the payload channels, channel manager 900 needs an active upstream signal on each payload channel. However, at any given time, some payload channels may not be allocated and some allocated channels may not be active. Thus, these payload channels do not provide the necessary upstream signals to the head end to monitor the quality of the payload channels. To compensate for these idle and unallocated payload channels, channel manager 900 places these channels in loop back mode to monitor the quality. In this case, channel manager 900 sets up the payload channel, transmits data to the ISU on the payload channel and the ISU transmits back specified data on an associated upstream payload channel. Channel manager 900 monitors these channels at the head end to determine error rates for the channels. Thus, the unallocated or idle payload channel can be monitored for errors the same as with active channels. The goal of channel manager 900 is to have payload established on all of the payload channels at a given time. However, it may be acceptable to monitor the performance of each channel at least once an hour if not active.

Channel manager 900 randomly selects and uses ISUs to monitor payload channels in loopback mode described above. This provides several benefits to the system. First, this allows channel manager to handle the diverse layout of a cable plant. Channel manager 900 sets up and uses paths over different legs from the various ODNs of the system. Further, random cycling of the ISUs used in the loop back mode allows the system to properly distribute power in the coaxial network. Specifically, this random selection of ISUs for loopback mode applies to concentration type services.

As described below, some ISUs are powered down when not active. When the ISU is powered down, the upstream modem at the head end detects this condition and sends a specified signal to the CXMC so that channel manager does not use the ISU for loop back purposes. Thus, powered down ISUs do not produce unnecessary errors.

Figure 68 is a flow chart that illustrates a method for monitoring payload channels by channel manager 900. Channel manager 900 reads parity error registers of the CXMU 56 are read every 10 milliseconds. Generally, the error counts are used to update the channel quality database and determine which (if any) channels require reallocation. The database of channel manager 900 contains an ongoing record of each channel. An accumulator sums the errors with previously recorded errors to update the database. The database organizes the history of the channels in categories such as: current ISU assigned to the channel, start of monitoring, end of monitoring, total error, errors in last day, in last week, number of seconds since last error, severe errors in last day, in last week, and current service type, such as ISDN, assigned to the channel. When the channel is a regular (non-loop back) payload channel, channel manager 900 determines whether the performance statistics in the database are within service specific threshold. When the statistics unacceptably exceed the threshold, channel manager 900 reallocates the channel using a "make before break" procedure to reduce the disruption from reallocating the channel. Thus, channel manager 900 allocates the new payload channel for the connection before deallocating the current payload channel.

Two issues presented by periodic parity monitoring as described above must be addressed in order to estimate the bit error rate corresponding to the observed count of parity errors in a monitoring period to determine if a channel is corrupted. The first is the nature of parity itself. Accepted practice for data formats using block error detection assumes that an errored block represents one bit of error, even though the error actually represents a large number of data bits. Due to the nature of the data transport system, errors injected into modulated data are expected to

randomize the data. This means that the average errored frame will consist of four (4) errored data bits (excluding the ninth bit). Since parity detects only odd bit errors, half of all errored frames are not detected by parity. Therefore, each parity (frame) error induced by transport interference represents an average of 8 (data) bits of error. Second, each monitoring parity error represents 80 frames of data (10 ms/125  $\mu$ s). Since the parity error is latched, all errors will be detected, but multiple errors will be detected as one error.

The bit error rate (BER) used as a basis for determining when to reallocate a channel has been chosen as  $10^{-3}$ . Therefore, the acceptable number of parity errors in a one second interval that do not exceed  $10^{-3}$  must be determined. To establish the acceptable parity errors, the probable number of frame errors represented by each observed (monitored) parity error must be predicted. Given the number of monitored parity errors, the probable number of frame errors per monitored parity error, and the number of bit errors represented by a frame (parity) error, a probable bit error rate can be derived.

A statistical technique is used and the following assumptions are made:

1. Errors have a Poisson distribution, and
2. If the number of monitored parity errors is small ( $< 10$ ) with respect to the total number of "samples" (100), the monitored parity error rate (MPER) reflects the mean frame error rate (FER).

Since a monitored parity error (MPE) represents 80 frames, assumption 2 implies that the number of frame errors (FEs) "behind" each parity error is equal to 80 MPER. That is, for 100 parity samples at 10 ms per sample, the mean number of frame errors per parity error is equal to 0.8 times the count of MPEs in one second. For example, if 3 MPEs are observed in a one second period, the mean number of FEs for each MPE is 2.4. Multiplying the desired bit error rate times the sample size and dividing by the bit errors per frame error yields the equivalent number of frame errors in the sample. The number of FEs is also equal to the product of the



number of MPEs and the number of FEs per MPE. Given the desired BER, a solution set for the following equation can be determined.

$$(MPE \frac{FE}{MPE}) = 0.8 MPE$$

The Poisson distribution, as follows, is used to compute the probability of a given number of FEs represented by a MPE ( $\chi$ ), and assumption 2, above, is used to arrive at the mean number of FEs per MPE ( $\mu$ ).

$$P(x) = \frac{e^{-\mu} \mu^x}{x!}$$

Since the desired bit error rate is a maximum, the Poisson equation is applied successively with values for  $\chi$  of 0 up to the maximum number. The sum of these probabilities is the probability that no more than  $\chi$  frame errors occurred for each monitored parity error.

The results for a bit error rate of  $10^{-3}$  and bit errors per frame error of 1 and 8 are shown in Table 11.

Table 11: Bit Error Rate Probability

Bit Errors per Frame Error	Monitored Parity Errors	Maximum Frame Errors/Monitored Parity Error ( $\chi$ )	Average Frame Errors/Monitored Parity Error ( $\mu$ )	Probability of BER $< 10^{-3}$
8	2	4	1.6	98%
	3	3	2.4	78%
	4	2	3.2	38%
1	8	8	6.4	80%
	9	7	7.2	56%
	10	7	8.0	45%

Using this technique, a value of 4 monitored parity errors detected during a

5

## 10

15

20

## 25

**The IDL is a standard payload channel that is dynamically assigned to**

transmit control data between HDT 12 and ISU 100 when the amount of data exceeds the parameters of the lower bandwidth IOC channel. The IDL channel can provide full duplex communication or simplex broadcast from HDT 12 to one or more ISUs 100. Channel manager 900 dynamically allocates the IDL channel as needed for non-time critical transport of data as described above.

The IDL messages in both directions are variable in length. The IDL data is transmitted over HFC distribution network 11 at a rate of 64 Kbps which is one byte per 8 kHz frame. The IDL channel uses one of the 240 payload channels and the protocol for transmitting IDL messages is handled by the processor on CXMU 56. The processor uses cyclical redundancy codes (CRC) and positive acknowledgments to manage error checking of IDL messages.

The IDL channel can be used to transmit various kinds of data. For example, the IDL channel can be used to provide data to an ISU to configure a payload channel for use with a specific protocol. For example, the IDL channel can be used to down load data to configure a payload channel for use with the LAPB protocol or any other appropriate protocol, including proprietary protocols. Similarly, the IDL channel can be used to download software to an ISU.

Transmission over the IDL channel has a lower priority than regular payload transmissions. Thus, channel manager 900 deallocates an IDL channel before completion of the data transmission when channel manager 900 receives a request that requires use of the payload channel that is currently allocated to the IDL.

Figure 69 is a flow chart that illustrates an embodiment of a method for allocating a payload channel to the ISU data link. At block 330a, channel manager 900 receives a request for an IDL channel. At block 332a, channel manager 900 determines whether a payload channel is available. If a payload channel is available, channel manager 900 allocates the payload channel to the IDL channel and the data is transmitted to the identified ISU. If, however, a channel is not available, channel manager determines whether one of the allocated channels is idle by checking the hook state of a line of a channel unit. If the line is on hook, then

channel manager 900 reallocates the channel to the IDL channel until the IDL transmission is complete. If however, channel manager receives a request for a communication link to the line of the channel unit, channel manager interrupts the IDL channel and reallocates the payload channel to the channel unit.

#### Power Down

Channel manager 900 can power down an ISU during periods of non-use to reduce energy costs of communication system 10. To power down the ISU, channel manager 900 must determine that all conditions for powering down the ISU are met. For example, channel manager 900 can determine if the lines of the channel units of the ISU provide service that can be powered down. Such services may include, for example, analog services such as POTS and COIN. Further, the lines must be idle. For example, channel manager 900 can determine if a line is idle based on the line's hook status or other appropriate criteria. Channel manager 900 checks the status of the lines each time a line goes from off-hook to on-hook. Channel manager 900 further checks the status of the lines every second to monitor the hook status. It is noted, however, that channel manager 900 will not power down an ISU that is the monitoring ISU for a subband.

When channel manager 900 determines that an ISU can be powered down, the ISU's transmitter is disabled. Head end 32 detects the loss of power to the ISU and sends an idle pattern upstream to the switch. An IOC control message to or from the IOC will power-up the ISU. The IOC traffic to or from the ISU indicates to the background task in charge of powering down ISUs to check the ISU against the criteria for powering down again.

It is understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reviewing the above description. The scope of the invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

Further Embodiment of Telephony Transport System

With reference to Figure 116, a general description of a hybrid fiber/coax communications network 1006 in accordance with the present invention shall be described. Telephony and video information from existing telephone and video services generally shown by trunk line 1008 is received by and processed by head end 1010. Head end 1010 includes a plurality of host distribution terminals (HDT) 1012 for telephony data interface and video host distribution terminal (VHDT) 1014 for video data interface. Host distribution terminals 1012 and VHDT 1014 include transmitters and receivers for communicating the video and telephony information between the video and telephony signal distribution network 1006 in accordance with the present invention and the existing telephony and video services as represented generally by trunk line 1008.

The video information is optically transmitted downstream via optical fiber line 1017 to splitter 1018 which splits the optical video signals for transmission on a plurality of optical fibers 1022 to a plurality of optical distribution nodes 1026. The HDT 1012 transmits optical telephony signals via optical fiber link 1020 to the optical distribution nodes 1026. The optical distribution nodes 1026 convert the optical video signals and telephony signals for transmission as electrical outputs via a coaxial distribution system 1007 to a plurality of remote units 1042. The electrical downstream video and telephony signals are distributed via a plurality of coaxial lines 1029 and coaxial taps 1034 of the coaxial distribution system 1007.

The remote units 1042 include means for transmitting upstream electrical data signals including telephony information from telephones 1076 and data terminals 1073 and in addition may include means for transmitting set top box information from set top boxes 1078. The upstream electrical data signals are provided by a plurality of remote units 1042 to an optical distribution node 1026 connected thereto. The optical distribution node 1026 converts the upstream electrical data signals to an upstream optical data signal for transmission via optical

fiber link 1020 to the head end 1010.

The present invention shall now be described in further detail with reference to Figures 116-123. The first part of the description shall primarily deal with downstream transmission and the second part of the description shall primarily be with regard to upstream transmission. The video and telephony distribution network 1006 in accordance with the present invention, includes head end 1010 which receives video and telephony information from video and telephony service providers via trunk line 1008. Head end 1010 includes a plurality of host distribution terminals 1012 and a video host distribution terminal 1014. The HDT 1012 includes a transmitters and receivers for communicating telephony information, such as T1, ISDN, or other data services information, to and from telephony service providers via trunk line 1008 and the VHDT 1014 includes a transmitters and receivers for communicating video information, such as cable TV video information and interactive data of subscribers to and from video service providers via trunk line 1008.

The VHDT 1014 transmits downstream optical signals to a splitter 1018 via video feeder optical fiber line 1017. The passive optical splitter 1018 effectively makes four copies of the downstream high bandwidth optical video signals. The duplicated downstream optical video signals are distributed to the correspondingly connected optical distribution nodes 1026. One skilled in the art will readily recognize that although four copies of the downstream video signals are created, that any number of copies may be made by an appropriate splitter and that the present invention is not limited to any specific number.

The splitter 1018 is a passive means for splitting broad band optical signals without the need to employ expensive broad band optical to electrical conversion hardware. Optical signal splitters are commonly known to one skilled in the art and available from numerous fiber optic component manufacturers such as Gould, Inc. In the alternative, active splitters may also be utilized. In addition, a cascaded chain of passive or active splitters would further multiply the number of

5 duplicated optical signals for application to an additional number of optical distribution nodes and therefore increase further the remote units serviceable by a single head end. Such alternatives are contemplated in accordance with the present invention as described by the accompanying claims.

10 The VHDT 1014 can be located in a central office, cable TV head end, or a remote site and broadcast up to about 112 NTSC channels. The VHDT 1014 includes a transmission system like that of a LiteAMp™ system available from American Lightwave Systems, Inc., currently a subsidiary of the assignee hereof. Video signals are transmitted optically by amplitude modulation of a 1300 nanometer laser source at the same frequency at which the signals are received (i.e. the optical transmission is a terahertz optical carrier which is modulated with the RF video signals). The downstream video transmission bandwidth is about 54-725 MHz. One advantage in using the same frequency for optical transmission of the video signal as the frequency of the video signals when received is to provide high bandwidth transmission with reduced conversion expense. This same-frequency transmission approach means that the modulation downstream requires optical to electrical conversion or proportional conversion with a photodiode and perhaps amplification, but no frequency conversion. In addition, there is no sample data bandwidth reduction and little loss of resolution.

15 Alternative embodiments of the VHDT may employ other modulation and mixing schemes or techniques to shift the video signals in frequency, and other encoding methods to transmit the information in a coded format. Such techniques and schemes for transmitting analog video data, in addition to those transmitting digital video data, are known to one skilled in the art and are contemplated in accordance with the spirit and scope of the present invention as described in the accompanying claims.

20 Telephony information is transmitted downstream by HDT 1012 via optical fiber link 1020 to a corresponding optical distribution node 1026. A more detailed block diagram of one of the HDTs 1012 is shown in Figure 117.

Each HDT 1012 includes an RF modem bank 1050 which receives telephony information via trunk line 1008. The RF modem bank 1050 includes four RF modem modules 1052 and a protection modem module 1054. Each RF modem module receives telephony information, for example time division multiplexed channel signals from a public switched telephone service, via trunk line 1008 and the telephony information modulates an analog carrier for transmission of the downstream optical telephony data by downstream optical telephony transmitter 1080 of downstream telephony electrical to optical converter 1064 to a corresponding distribution node 1026. Each RF modem module includes a transceiver 1053 and provides a downstream electrical telephony signal in one of four frequency bandwidths, each bandwidth being about 6 MHz in width like that of a CATV channel. Each 6 MHz bandwidth channel transmits data at 22 Mbits/sec and can provide for transmission of 8T1 digital telephone signals; T1 being a conventional telephone signal where 24 voice channels are sampled at an 8 kHz rate, with 8 bits per sample (each 8 bit conversation sample is termed a DS0). Each of these signals from the four RF modem modules 1052 are transmitted via coax patch cables to a combiner 82 of downstream telephony electrical to optical converter 1064 for transmission by optical transmitter 1080. Therefore, the spectrum for the downstream optical telephony data is four separated 6 MHz frequency bands containing 22 Mbits/sec of data within each 6 MHz bandwidth. The four 6 MHz frequency bands, separated by a guard band as is known to one skilled in the art, are transmitted in about the 725-800 MHz bandwidth.

Any number of modulation techniques may be used for transmission of the telephony information downstream. The transmission downstream is point to multipoint transmission using broadcast type transmission schemes. The modulation techniques utilized and performed by RF modem module 1052 may include quadrature phase shift keying (QPSK), quadrature amplitude modulation (QAM), or other modulation techniques for providing the desired data rate. Modulation techniques, such as QPSK and QAM, are known to those skilled



in the art and the present invention contemplates the use of any such modulation techniques for downstream broadcast transmission.

The electrical to optical converter 1064 includes two transmitters 1080 for downstream telephony transmission to protect the telephony data transmitted. These transmitters are conventional and relatively inexpensive narrow band laser transmitters. One transmitter is in standby if the other is functioning properly. Upon detection of a fault in the operating transmitter, controller 1060 switches transmission to the standby transmitter. In contrast, the transmitter of the VHDT 1014 is relatively expensive as compared to the transmitters of HDT 1012 as it is a broad band analog DFB laser transmitter. Therefore, protection of the video information, non-essential services unlike telephony data, is left unprotected. By splitting the telephony data transmission from the video data transmission, protection for the telephony data alone can be achieved. If the video data information and the telephony data were transmitted over one optical fiber line by an expensive broad band analog laser, economies may dictate that protection for telephony services may not be possible. Therefore, separation of such transmission is of importance.

As an alternative embodiment for providing transmission of optical video and telephony signals to the optical distribution nodes 1027 from head end 1010 as shown in Figure 120, the HDT 1012 and VHDT 1014 can utilize the same optical transmitter and the same optical fiber line 1016. The signal then is split by splitter 1018 and four split signals are provided to the optical distribution nodes 1027 for distribution to the remote units 1042 by the coaxial distribution system 1007 as further discussed below. However, as described above, the optical transmitter utilized would be relatively expensive due to its broad band capabilities, lessening the probabilities of being able to afford protection for essential telephony services.

As one skilled in the art will recognize, optical link 1020, as shown in Figure 117, may include four fibers, two for transmission downstream

from electrical to optical converter 1064 and two for transmission upstream to optical to electrical converter 1066. With the use of directional couplers, the number of such fibers may be cut in half. In addition, the number of protection transmitters and fibers utilized may vary as known to one skilled in the art and any listed number is not limiting to the present invention as described in the accompanying claims.

RF modem bank 1050 includes a protection RF modem module 1054 with a transceiver 1053 connected to combiner 1082 of electrical to optical converter 1064. Protection RF modem module 1054 is further coupled to controller 1060. When a fault is detected with regard to the transmission of one of the RF modem modules 1052, a signal is generated and applied to an input 1062 of controller 1060. Controller 1060 is alerted to the fault and provides appropriate signaling to switch the protection RF modem module 1054 for the faulted RF modem such that the protection RF modem module 1054 transmits within the 6 MHz bandwidth of the faulted RF modem module 1052 so that the four 6 MHz bandwidth signal transmission is continued on optical fiber link 1020. The use of one protection RF modem module 1054 for four RF modem modules 1052 is only one embodiment of the present invention and the number of protection RF modem modules relative to RF modem modules may vary as known to one skilled in the art and described in the accompanying claims. As shown in Figure 123, RF modem bank 1050 may include one protection module 1054 for each RF modem module 1052. In this embodiment, the RF modem bank 1050 includes three RF modem modules 1052 and three protection modules 1054 for one-to-one protection.

An optical distribution node 1026 as shown in Figure 118 receives both the downstream optical telephony signal and the split downstream optical video signal. The downstream optical video signal is applied by the optical fiber 1022 from splitter 1018 to a downstream video receiver 1120 of optical distribution node 1026. The optical distribution node 1026 further includes downstream telephony receiver 1121 for receiving the downstream optical telephony

signal on optical link 1020. The optical video receiver 1120 utilized is like that available in the LiteAMP™ product line available from American Lightwave Systems, Inc. The converted signal from video receiver 1120, proportionally converted utilizing photodiodes, is applied to bridger amplifier 1127 along with the converted telephony signal from downstream telephony receiver 1121. The bridging amplifier 1127 simultaneously applies four downstream electrical telephony and video signals to diplex filters 1134. The diplex filters 1134 allow for full duplex operation by separating the transmit and receive functions when signals of two different frequency bandwidths are utilized for upstream and downstream transmission. There is no frequency conversion performed at the optical distribution nodes with respect to the video or downstream telephony signals as the signals are passed through the optical distribution nodes to the remote units via the coaxial distribution system in the same frequency bandwidth as they are received.

After the optical distribution node 1026 has received downstream optical video signals via optical link 1022 and downstream optical telephony signals via optical link 1020 and such signals are converted to downstream electrical video and telephony signals, the four outputs of the optical distribution nodes 1026 are applied to four coaxial cables 1029 of coaxial cable distribution system 1007 for transmission of the downstream electrical video and telephony signals to the remote units 1042; such transmission occurs in about the 725-800 MHz bandwidth for telephony signals and about the 54-725 MHz bandwidth for the downstream electrical video signals. Each optical distribution node 1026 provides for the transmission over a plurality of coaxial cables 1029 and any number of outputs is contemplated in accordance with the present invention as described in the accompanying claims.

As shown in Figure 116, each coaxial cable 1029 can provide a significant number of remote units with downstream electrical video and telephony signals through a plurality of coaxial taps 1034. Coaxial taps are commonly known to one skilled in the art and act as passive bidirectional pick-offs of electrical

5 The downstream electrical video and telephony signals are provided from the coaxial taps to the remote units 1042 in a number of different ways. In one embodiment, the signal from the coaxial tap 1034 is provided to a home integrated service unit 1070 as shown in Figure 119. The home integrated service unit 1070 of Figure 119 includes a power tap 1099 coupled to a conventional power supply and ring generator 1101. The downstream electrical video and telephony signals are provided to a tap 1097 for application of the signals to both duplex filter 1110 and ingress filter 1098. The downstream video signal is provided from ingress filter 1098 to video equipment 1072 via set top box 1078. The downstream telephony signal is applied from duplex filter 1110 to RF demodulator 1104 of RF modem module 1102 and the demodulated signal is applied to an applicable service interface for processing and connection to user equipment. For example, the RF demodulated signal is processed via Plain Old Telephone Service (POTS) service interface 1112 for output on twisted pairs 1118 to telephone 1076 by POTS connection 1114. The other service interfaces such as ISDN interface or a T1 interface perform their conventional functions as are known to those skilled in the art for transmittal of such information on outputs thereof to user equipment.

25

**The set top box 1078 is an optional element in the network.**

1006. It may include an additional modem for sending interactive data therefrom back to head end 1010 at frequencies unused by the video and telephony transmissions. Upstream transmission of such data is further discussed below.

Depending on the modulation processing techniques utilized at the head end 1010, the RF demodulator 1104 would include circuitry capable of demodulating the modulated signal. For example, if QPSK modulation is utilized then the demodulator would include processing circuitry capable of demodulating a QPSK modulated waveform as is known to one skilled in the art.

In another embodiment of providing downstream electrical video and telephony signals from the coaxial taps 1034 to remote units 1042, as shown in Figure 116, a separate coaxial line from coaxial tap 1034 is utilized to provide transmission of the signals therefrom to set top box 1078, and thus for providing the downstream video signals to video equipment unit 1072. In such a case, a second coaxial line from coaxial tap 1034 would be utilized to provide the downstream telephony signals to a multiple integrated service unit (MISU) 1044 which would be much like the home integrated service unit 1070 as described with regard to Figure 119 except lacking an ingress filter 1098 and tap 1097. Unlike home integrated service unit 1070, the MISU 1044 would be utilized to service several remote units 1042 with telephony services via various service interfaces. Whether the video and telephony signals are provided to the curb with use of the MISU 1044 or whether the video and telephony signals are provided directly to a home integrated service unit is strictly one of application and either can be utilized with regard to the same or different coaxial taps 1034 and within the same or different coaxial distribution systems 1007.

In addition, an optional network interface device (NID) 1075 is utilized in the connection of telephone services to the remote units 1042, whether they are homes or businesses, as is known to those skilled in the art and as shown in Figure 116. The NID is generally shown by block 1070 representing the home integrated service unit but is not shown in the detail of Figure 119. The NID

performs service functions for the telephone service provider such as looping back signals to the service provider that reach the NID so as to indicate whether a failure has occurred somewhere in transmission to the NID or in connections from the NID to the user equipment when a failure is reported to the service provider.

5           The above description primarily involves the downstream transmission of video and telephony information from head end 1010 to remote units 1042. The upstream transmission of interactive data from set top boxes 1078 and other data, for example telephony from telephones 1076, shall now be described with reference to Figures 116-123. The description shall be limited to transmission  
10       from remote units via home integrated service units as transmission from an MISU is substantially similar and easily ascertainable from such description. Home integrated service unit 1074 provides set top box information from set top box 1078 and telephony information from the service interfaces 1112, including information  
15       from telephone 1076, to the optical distribution node 1026 connected thereto by the same coaxial path as for the downstream communication. The set top box signals are transmitted by a separate RF modem of the video service provider at a relatively low frequency in the bandwidth of about 5 to 40 MHz which is unused by telephony and video services. The telephony signals are also transmitted upstream in the 5-40  
20       MHz bandwidth, usually from 10 MHz to 30 MHz. This 5-40 MHz bandwidth is reused in each coaxial path 1029 from each remote unit 1042 to the respectively connected optical distribution node 1026. As such, upstream electrical telephony data signals from the remote units are transmitted at the same reused frequency bandwidth of 5-40 MHz on each coaxial line 1029 for input to the optical  
25       distribution node 1026. Therefore, as shown in Figure 118, four upstream electrical telephony signals, each in the 5-40 MHz bandwidth, are input to optical distribution node 1026, via the respectively connected coaxial cables 1029.

          The upstream transmission from an integrated service unit for multipoint to point transmission utilizes time multiplexing techniques, although any of a number of multiple access techniques known to those skilled in the art are

contemplated in accordance with the present invention. All the remote units are designated a time slot for transmission. In such a case each remote unit must transmit at a particular time to maintain multiple access with the timing being supplied using data on the downstream paths. The upstream data is transmitted on a bit-by-bit basis. With each remote unit assigned a time slot, the RF modem 1102 of the unit knows that it will not interfere with the others because it has determined the time delay for each one of them and each RF modem 1102 is signaled to transmit at a precise time. Due to the high volumes of multiplexed serial data from several outlining remote stations and limited bandwidth for transmission, short pulse durations are required for better resolution of the data transmitted to the head end 1010. Although the data modulates a carrier and is transmitted in the 5 to 40 MHz bandwidth by RF modulator 1108, because of the limited bandwidth in the upstream direction, a pulse shaping network at each remote unit is used to generate raised cosine pulses for the rectangular or square wave bit-by-bit stream of data transmitted along the coaxial cable in the coaxial network.

An optimal pulse shape for transmission in a band limited coaxial cable network is determined by the use of Fourier calculations with a given set of boundary conditions. Also, the Fourier calculations implement a spectral limitation constraint for the purposes of limiting the spectral content of the optimal pulse shape. Limiting the spectral content of the pulse shape serves two functions. The first function is to limit the spectral characteristics of the optimal pulse shape in order to prevent phase dispersion at the receiving end of the transmission system. The second benefit from the spectral limitation constraint is to allow the use of relatively simple finite impulse response filters with a minimal number of taps.

In one embodiment of the pulse shaping network as shown in Figure 121, 50 nanosecond pulses from the RF modulator 1108 of RF modem 1102 are transmitted to a pulse sequencer 1301 for uniform digitization. The output from the pulse sequencer is then applied to a ten tapped finite impulse response filter (FIR filter) 1302 with associated electronics 1303 to provide the addition and subtraction

necessary for the filtering process. The output is sent to a line driver circuit for output to the coaxial cable through duplex filter 1110. The optimal pulse waveform is a raised cosine waveform. Using such pulse shaping techniques, overcomes the difficulty of sending extremely short pulse duration information along a band limited coaxial cable.

The upstream electrical telephony signals from a plurality of remote units, including signals from the RF modems 1102 and from modems in set top boxes 1078, are transmitted to the respectively connected optical distribution node 1026 as shown in Figure 118 via the individual coaxial cables 1029. The upstream electrical signals are applied to a duplex filter 1134 respectively connected to a coaxial cable 1029. One of the duplex filters 1134 passes the upstream electrical telephony signal applied thereto through to combiner 1125 while the other duplex filters pass the upstream electrical telephony signals applied thereto to frequency shifters 1128, 1130, and 1132. Frequency shifter 1128 shifts the upstream electrical telephony signal into the 50-85 MHz bandwidth, frequency shifter 1130 shifts another upstream electrical telephony signal into the 100-135 MHz bandwidth and frequency shifter 1132 shifts the other upstream electrical telephony signal into the 150-185 MHz bandwidth. The shifted signals are combined by combiner 1125 and provided to upstream telephony and set top control transmitters 1123. The conventional optical transmitters 1123 transmit the upstream electrical telephony signal as an upstream optical telephony signal to head end 1010 via fiber optic link 1020. Once again, two transmitters are available for transmission, one in standby mode, like that in the downstream transmission path.

The upstream optical telephony signals are received by upstream telephony and set top box receiver 1084 of optical to electrical converter block 1066. The upstream optical telephony signals are converted, split, and all the split electrical signals in the 5-40 MHz, 50-85 MHz, 100-135 MHz, and 150-185 MHz are frequency shifted back to the 5-40 MHz bandwidth by frequency shifters 1086, 1088, and 1090 with the exception of the signal already in the 5-40 MHz



bandwidth which is passed through with the other frequency shifted signals from the frequency shifters to RF switch 1094. A combined signal in the 5-40 MHz bandwidth from combiner 1092 is provided to the VHDT and the signal is processed for obtaining the interactive information transmitted from set top boxes 1078. The RF switch 1094 is controlled by controller 1060 and provides the upstream telephony signals to the transceivers 1053 of the corresponding RF modems 1052. The upstream telephony signals are then demodulated by RF modem modules 1052 and the telephony data is provided to the service providers via trunk line 1008. The RF modem modules 1052 include RF demodulator corresponding to the modulation techniques utilized to transmit the information upstream so such information can be recovered.

As discussed previously, the controller 1060 switches protection RF modem module 1054 for a transmitting RF modem module 1052 in the downstream communication when a fault is detected in that module. The controller also provides signaling for switching the RF switch 1094 such that the information which would have been provided to the faulted RF modem module 1052 is applied to the transceiver of the protection RF modem module 1054. Therefore, the protection modem module 1054 is then fully within the transmit and receive loop of the system.

As shown in Figure 122, an alternative embodiment of the present invention includes an optical to electrical converter 1066 wherein the received optical upstream telephony signal is converted by receivers 1084 and the entire upstream electrical signal in the 5-200 MHz bandwidths is applied to the transceivers 1053 of the RF modem modules 1052. The RF modem modules 1052 are then operated under control of controller 1060 which assigns the RF modem module a carrier frequency to tune to for the recovery of telephony information; the assigned frequency being a function of the frequency shifting of the upstream signal. The electrical signal is still separated and frequency shifted by frequency shifters 1086, 1088 and 1090 except for the signal already in the 5-40 MHz bandwidth and

then combined by combiner 1092 for application to VHDT 1014.

In this embodiment, the switching of the protection modem module 1054 into the system is accomplished through the controller 1060. When the controller 1060 detects and indicates a faulted modem module 1052, the controller 1060 assigns the frequency previously assigned to the faulted RF modem module to the protection module, thus establishing the protection RF modem module 1054 fully within the transmit and receive loop.

In another embodiment shown in Figure 123 including one-to-one protection for the RF modem module, neither the RF switch used for protection switching for the configuration of Figure 123 nor the additional control required for protection switching for the configuration of Figure 122 is necessary. In this embodiment, the same electrical signal provided to the RF modem modules 1052 is applied to the corresponding protection module 1054, thus only a control signal indicating which module is to be used for transmission or reception is required for the one-to-one protection.

It is to be understood, however, that even though numerous characteristics of the present invention have been set forth in the foregoing description, together with details of the structure and function of the invention, the disclosure is illustrative and changes in matters of shape, size, number, and arrangement of the elements may be made within the principles of the invention and to the full extent indicated by the broad general meaning of the terms in which the appending claims are expressed.